

Machine-Learning Research

Four Current Directions

Thomas G. Dietterich

■ Machine-learning research has been making great progress in many directions. This article summarizes four of these directions and discusses some current open problems. The four directions are (1) the improvement of classification accuracy by learning ensembles of classifiers, (2) methods for scaling up supervised learning algorithms, (3) reinforcement learning, and (4) the learning of complex stochastic models.

The last five years have seen an explosion in machine-learning research. This explosion has many causes: First, separate research communities in symbolic machine learning, computational learning theory, neural networks, statistics, and pattern recognition have discovered one another and begun to work together. Second, machine-learning techniques are being applied to new kinds of problem, including knowledge discovery in databases, language processing, robot control, and combinatorial optimization, as well as to more traditional problems such as speech recognition, face recognition, handwriting recognition, medical data analysis, and game playing.

In this article, I selected four topics within machine learning where there has been a lot of recent activity. The purpose of the article is to describe the results in these areas to a broader AI audience and to sketch some of the open research problems. The topic areas are (1) ensembles of classifiers, (2) methods for scaling up supervised learning algorithms, (3) reinforcement learning, and (4) the learning of complex stochastic models.

The reader should be cautioned that this article is not a comprehensive review of each of these topics. Rather, my goal is to provide a representative sample of the research in each of these four areas. In each of the areas, there are many other papers that describe relevant work. I apologize to those authors whose work I was unable to include in the article.

Ensembles of Classifiers

The first topic concerns methods for improving accuracy in supervised learning. I begin by introducing some notation. In supervised learning, a learning program is given training examples of the form $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)\}$ for some unknown function $y = f(\mathbf{x})$. The \mathbf{x}_i values are typically vectors of the form $\langle x_{i,1}, x_{i,2}, \dots, x_{i,n} \rangle$ whose components are discrete or real valued, such as height, weight, color, and age. These are also called the features of \mathbf{x}_i . I use the notation x_j to refer to the j th feature of \mathbf{x}_i . In some situations, I drop the i subscript when it is implied by the context.

The y values are typically drawn from a discrete set of classes $\{1, \dots, K\}$ in the case of classification or from the real line in the case of regression. In this article, I focus primarily on classification. The training examples might be corrupted by some random noise.

Given a set S of training examples, a learning algorithm outputs a classifier. The *classifier* is a hypothesis about the true function f . Given new \mathbf{x} values, it predicts the corresponding y values. I denote classifiers by h_1, \dots, h_l .

An *ensemble of classifiers* is a set of classifiers whose individual decisions are combined in some way (typically by weighted or unweighted voting) to classify new examples. One of the most active areas of research in supervised learning has been the study of methods for constructing good ensembles of classifiers. The main discovery is that ensembles are often much more accurate than the individual classifiers that make them up.

An ensemble can be more accurate than its component classifiers only if the individual classifiers disagree with one another (Hansen and Salamon 1990). To see why, imagine that we have an ensemble of three classifiers: $\{h_1, h_2, h_3\}$, and consider a new case \mathbf{x} . If the three clas-




Figure 1. The Probability That Exactly ℓ (of 21) Hypotheses Will Make an Error, Assuming Each Hypothesis Has an Error Rate of 0.3 and Makes Its Errors Independently of the Other Hypotheses.

sifiers are identical, then when $h_1(\mathbf{x})$ is wrong, $h_2(\mathbf{x})$ and $h_3(\mathbf{x})$ are also wrong. However, if the errors made by the classifiers are uncorrelated, then when $h_1(\mathbf{x})$ is wrong, $h_2(\mathbf{x})$ and $h_3(\mathbf{x})$ might be correct, so that a majority vote correctly classifies \mathbf{x} . More precisely, if the error rates of L hypotheses h_i are all equal to $p < 1/2$ and if the errors are independent, then the probability that the majority vote is wrong is the area under the binomial distribution where more than $L/2$ hypotheses are wrong. Figure 1 shows this area for a simulated ensemble of 21 hypotheses, each having an error rate of 0.3. The area under the curve for 11 or more hypotheses being simultaneously wrong is 0.026, which is much less than the error rate of the individual hypotheses.

Of course, if the individual hypotheses make uncorrelated errors at rates exceeding 0.5, then the error rate of the voted ensemble increases as a result of the voting. Hence, the key to successful ensemble methods is to construct individual classifiers with error rates below 0.5 whose errors are at least somewhat uncorrelated.

Methods for Constructing Ensembles

Many methods for constructing ensembles have been developed. Some methods are general, and they can be applied to any learning algorithm. Other methods are specific to partic-

ular algorithms. I begin by reviewing the general techniques.

Subsampling the Training Examples

The first method manipulates the training examples to generate multiple hypotheses. The learning algorithm is run several times, each time with a different subset of the training examples. This technique works especially well for *unstable learning algorithms*—algorithms whose output classifier undergoes major changes in response to small changes in the training data. Decision tree, neural network, and rule-learning algorithms are all unstable. Linear-regression, nearest-neighbor, and linear-threshold algorithms are generally stable.

The most straightforward way of manipulating the training set is called *bagging*. On each run, bagging presents the learning algorithm with a training set that consists of a sample of m training examples drawn randomly with replacement from the original training set of m items. Such a training set is called a *bootstrap replicate* of the original training set, and the technique is called *bootstrap aggregation* (Breiman 1996a). Each bootstrap replicate contains, on the average, 63.2 percent of the original training set, with several training examples appearing multiple times.

Another training-set sampling method is to construct the training sets by leaving out disjoint subsets of the training data. For example, the training set can be divided randomly into 10 disjoint subsets. Then, 10 overlapping training sets can be constructed by dropping out a different one of these 10 subsets. This same procedure is used to construct training sets for tenfold cross-validation; so, ensembles constructed in this way are sometimes called *cross-validated committees* (Parmanto, Munro, and Doyle 1996).

The third method for manipulating the training set is illustrated by the ADABOOST algorithm, developed by Freund and Schapire (1996, 1995) and shown in figure 2. Like bagging, ADABOOST manipulates the training examples to generate multiple hypotheses. ADABOOST maintains a probability distribution $p_i(\mathbf{x})$ over the training examples. In each iteration ℓ , it draws a training set of size m by sampling with replacement according to the probability distribution $p_i(\mathbf{x})$. The learning algorithm is then applied to produce a classifier h_ℓ . The error rate ϵ_ℓ of this classifier on the training examples (weighted according to $p_i(\mathbf{x})$) is computed and used to adjust the probability distribution on the training examples. (In figure 2, note that the probability distribution is obtained by normalizing a set of weights $w_i(i)$ over the training examples.)

The effect of the change in weights is to place more weight on training examples that were misclassified by h_ℓ and less weight on examples that were correctly classified. In subsequent iterations, therefore, ADABOOST constructs progressively more difficult learning problems.

The final classifier, h_f , is constructed by a weighted vote of the individual classifiers. Each classifier is weighted according to its accuracy for the distribution p_ℓ that it was trained on.

In line 4 of the ADABOOST algorithm (figure 2), the base learning algorithm *Learn* is called with the probability distribution p_ℓ . If the learning algorithm *Learn* can use this probability distribution directly, then this procedure generally gives better results. For example, Quinlan (1996) developed a version of the decision tree-learning program C4.5 that works with a weighted training sample. His experiments showed that it worked extremely well. One can also imagine versions of back propagation that scaled the computed output error for training example (\mathbf{x}, y) by the weight $p_\ell(\mathbf{x})$. Errors for important training examples would cause larger gradient-descent steps than errors for unimportant (low-weight) examples.

However, if the algorithm cannot use the probability distribution p_ℓ directly, then a training sample can be constructed by drawing a random sample with replacement in proportion to the probabilities p_ℓ . This procedure makes ADABOOST more stochastic, but experiments have shown that it is still effective.

Figure 3 compares the performance of C4.5 to C4.5 with ADABOOST.M1 (using random sampling). One point is plotted for each of 27 test domains taken from the Irvine repository of machine-learning databases (Merz and Murphy 1996). We can see that most points lie above the line $y = x$, which indicates that the error rate of ADABOOST is less than the error rate of C4.5. Figure 4 compares the performance of bagging (with C4.5) to C4.5 alone. Again, we see that bagging produces sizable reductions in the error rate of C4.5 for many problems. Finally, figure 5 compares bagging with boosting (both using C4.5 as the underlying algorithm). The results show that the two techniques are comparable, although boosting appears to still have an advantage over bagging.

Manipulating the Input Features A second general technique for generating multiple classifiers is to manipulate the set of input features available to the learning algorithm. For example, in a project to identify volcanoes on Venus, Cherkauer (1996) trained an ensemble of 32 neural networks. The 32 networks were

```

Input: a set  $S$ , of  $m$  labeled examples:  $S = \{(\mathbf{x}_i, y_i), i = 1, 2, \dots, m\}$ ,
labels  $y_i \in Y = \{1, \dots, K\}$ 
LEARN (a learning algorithm)
a constant  $L$ .

[1] initialize for all  $i$ :  $w_1(i) := 1/m$            initialize the weights
[2] for  $\ell = 1$  to  $L$  do
[3]   for all  $i$ :  $p_\ell(i) := w_\ell(i) / (\sum_i w_\ell(i))$    compute normalized weights
[4]    $h_\ell := \text{LEARN}(p_\ell)$                        call LEARN with normalized weights.
[5]    $\epsilon_{\eta\ell} := \sum_i p_\ell(i) [h_\ell(\mathbf{x}_i) \neq y_i]$    calculate the error of  $h_\ell$ 
[7]   if  $\epsilon_{\eta\ell} > 1/2$  then
[8]      $L := \ell - 1$ 
[9]     goto 13
[10]   $\beta_{\eta\ell} := \epsilon_{\eta\ell} / (1 - \epsilon_{\eta\ell})$ 
[11]  for all  $i$ :  $w_{\ell+1}(i) := w_\ell(i) \beta_{\eta\ell}^{-1} [h_\ell(\mathbf{x}_i) \neq y_i]$    compute new weights
[12] end for

[13] Output:  $h_f(\mathbf{x}) = \underset{y \in Y}{\operatorname{argmax}} \sum_{\ell=1}^L \left( \log \frac{1}{\beta_{\eta\ell}} \right) [h_\ell(\mathbf{x}) = y]$ 

```

Figure 2. The ADABOOST.M1 Algorithm.

The formula $[E]$ is 1 if E is true and 0 otherwise.




Figure 3. Comparison of ADABOOST.M1 (applied to C4.5) with C4.5 by Itself.

Each point represents 1 of 27 test domains. Points lying above the diagonal line exhibit lower error with ADABOOST.M1 than with C4.5 alone. Based on data from Freund and Schapire (1996). As many as 100 hypotheses were constructed.

based on 8 different subsets of the 119 available input features and 4 different network sizes. The input-feature subsets were selected (by hand) to group features that were based on different image-processing operations (such as principal component analysis and the fast Fourier transform). The resulting ensemble classifier was able to match the performance of human experts in identifying volcanoes.




Figure 4. Comparison of Bagging (applied to C4.5) with C4.5 by Itself.

Each point represents 1 of 27 test domains. Points lying above the diagonal line exhibit lower error with bagging than with C4.5 alone. Based on data from Freund and Schapire (1996). Bagging voted 100 classifiers.




Figure 5. Comparison of Bagging (applied to C4.5) with ADABOOST.M1 (applied to C4.5).

Each point represents 1 of 27 test domains. Points lying above the diagonal line exhibit lower error with boosting than with bagging. Based on data from Freund and Schapire (1996).

Tumer and Ghosh (1996) applied a similar technique to a sonar data set with 25 input features. However, they found that deleting even a few of the input features hurt the performance of the individual classifiers so much that the voted ensemble did not perform well. Obviously, this technique only works when

the input features are highly redundant.

Manipulating the Output Targets A third general technique for constructing a good ensemble of classifiers is to manipulate the y values that are given to the learning algorithm. Dietterich and Bakiri (1995) describe a technique called *error-correcting output coding* (ECOC). Suppose that the number of classes, K , is large. Then, new learning problems can be constructed by randomly partitioning the K classes into two subsets: A_ℓ and B_ℓ . The input data can then be relabeled so that any of the original classes in set A_ℓ are given the derived label 0, and the original classes in set B_ℓ are given the derived label 1. These relabeled data are then given to the learning algorithm, which constructs a classifier h_ℓ . By repeating this process L times (generating different subsets A_ℓ and B_ℓ), we obtain an ensemble of L classifiers h_1, \dots, h_L .

Now, given a new data point \mathbf{x} , how should we classify it? The answer is to have each h_ℓ classify \mathbf{x} . If $h_\ell(\mathbf{x}) = 0$, then each class in A_ℓ receives a vote. If $h_\ell(\mathbf{x}) = 1$, then each class in B_ℓ receives a vote. After each of the L classifiers has voted, the class with the highest number of votes is selected as the *prediction* of the ensemble.

An equivalent way of thinking about this method is that each class j is encoded as an L -bit code word C_j , where bit ℓ is 1 if and only if $j \in B_\ell$. The ℓ -th learned classifier attempts to predict bit ℓ of these code words. When the L classifiers are applied to classify a new point \mathbf{x} , their predictions are combined into an L -bit string. We then choose the class j whose code word C_j is closest (in Hamming distance) to the L -bit output string. Methods for designing good error-correcting codes can be applied to choose the code words C_j (or, equivalently, subsets A_ℓ and B_ℓ).

Dietterich and Bakiri report that this technique improves the performance of both the C4.5 and back-propagation algorithms on a variety of difficult classification problems. Recently, Schapire (1997) showed how ADABOOST can be combined with error-correcting output coding to yield an excellent ensemble-classification method that he calls ADABOOST.OC. The performance of the method is superior to the ECOC method (and bagging) but essentially the same as another (complex) algorithm called ADABOOST.M2. Hence, the main advantage of ADABOOST.OC is implementation simplicity: It can work with any learning algorithm for solving two-class problems.

Ricci and Aha (1997) applied a method that combines error-correcting output coding with feature selection. When learning each classi-

Task	Test-Set Size	c4.5	200-Fold Bootstrap c4.5	200-Fold Random c4.5
Vowel	462	0.5758	0.5152	0.4870 ¹
Soybean	376	0.1090	0.0984	0.1090
Part of Speech	3060	0.0827	0.0765	0.0788 ²
NETTALK	7242	0.3000	0.2670 ³	0.2500 ³
Letter Recognition	4000	0.2010	0.0038 ³	0.0000 ³

1. Difference from c4.5 significant at $p < 0.05$.
2. 256-fold random.
3. Difference from c4.5 significant at $p < 0.001$.

Table 1. Results on Five Domains (best error rate in boldface).

er, h_i , they apply feature-selection techniques to choose the best features for learning the classifier. They obtained improvements in 7 of 10 tasks with this approach.

Injecting Randomness The last general-purpose method for generating ensembles of classifiers is to inject randomness into the learning algorithm. In the back-propagation algorithm for training neural networks, the initial weights of the network are set randomly. If the algorithm is applied to the same training examples but with different initial weights, the resulting classifier can be different (Kolen and Pollack 1991).

Although this method is perhaps the most common way of generating ensembles of neural networks, manipulating the training set might be more effective. A study by Parmanto, Munro, and Doyle (1996) compared this technique to bagging and tenfold cross-validated committees. They found that cross-validated committees worked best, bagging second best, and multiple random initial weights third best on one synthetic data set and two medical diagnosis data sets.

For the c4.5 decision tree algorithm, it is also easy to inject randomness (Kwok and Carter 1990). The key decision of c4.5 is to choose a feature to test at each internal node in the decision tree. At each internal node, c4.5 applies a criterion known as the *information-gain ratio* to rank the various possible feature tests. It then chooses the top-ranked feature-value test. For discrete-valued features with V values, the decision tree splits the data into V subsets, depending on the value of the chosen feature. For real-valued features, the decision tree splits the data into two subsets, depending on whether the value of the chosen feature is above or below a chosen threshold. Dietterich and Kong (1995) implemented a variant of c4.5 that chooses randomly (with equal probability) among the top 20 best tests. Table 1 compares a single run of c4.5 to ensembles of

200 classifiers constructed by bagging c4.5 and injecting randomness into c4.5. The results show that injecting randomness obtains the best performance in three of the domains. In particular, notice that injected randomness obtains perfect test-set performance in the letter-recognition task.

Ali and Pazzani (1996) injected randomness into the FOIL algorithm for learning Prolog-style rules. FOIL works somewhat like c4.5 in that it ranks possible conditions to add to a rule using an information-gain criterion. Ali and Pazzani computed all candidate conditions that scored within 80 percent of the top-ranked candidate and then applied a weighted random-choice algorithm to choose among them. They compared ensembles of 11 classifiers to a single run of FOIL and found statistically significant improvements in 15 of 29 tasks and statistically significant loss of performance in only one task. They obtained similar results using 11-fold cross-validation to construct the training sets.

Raviv and Intrator (1996) combine bootstrap sampling of the training data with injecting noise into the input features for the learning algorithm. To train each member of an ensemble of neural networks, they draw training examples with replacement from the original training data. The \mathbf{x} values of each training example are perturbed by adding Gaussian noise to the input features. They report large improvements in a synthetic benchmark task and a medical diagnosis task.

A method closely related to these techniques for injecting randomness is the *Markov chain Monte Carlo (MCMC) method*, which has been applied to neural networks by MacKay (1992) and Neal (1993) and to decision trees by Chipman, George, and McCulloch (1996). The basic idea of the MCMC method (and related methods) is to construct a Markov process that generates an infinite sequence of hypotheses h_i . In a Bayesian setting, the goal is to generate each

hypothesis h_t with probability $P(h_t|S)$, where S is the training sample. $P(h_t|S)$ is computed in the usual way as the (normalized) product of the likelihood $P(S|h_t)$ and the prior probability $P(h_t)$ of h_t . To apply MCMC, we define a set of operators that convert one h_t into another. For a neural network, such an operator might adjust one of the weights in the network. In a decision tree, the operator might interchange a parent and a child node in the tree or replace one node with another. The MCMC process works by maintaining a current hypothesis h_t . At each step, it selects an operator, applies it (to obtain h_{t+1}), and then computes the likelihood of the resulting classifier on the training data. It then decides whether to keep h_{t+1} or discard it and go back to h_t . Under various technical conditions, it is possible to prove that a process of this kind eventually converges to a stationary probability distribution in which the h_t 's are sampled in proportion to their posterior probabilities. In practice, it can be difficult to tell when this stationary distribution is reached. A standard approach is to run the Markov process for a long period (discarding all generated classifiers) and then collect a set of L classifiers from the Markov process. These classifiers are then combined by weighted vote according to their posterior probabilities.

Algorithm-Specific Methods for Generating Ensembles In addition to these general-purpose methods for generating diverse ensembles of classifiers, there are several techniques that can be applied to the back-propagation algorithm for training neural networks.

Rosen (1996) trains several neural networks simultaneously and forces the networks to be diverse by adding a correlation penalty to the error function that back propagation minimizes. Specifically, during training, Rosen keeps track of the correlations in the predictions of each network. He applies back propagation to minimize an error function that is a sum of the usual squared output prediction error and a term that measures the correlations with the other networks. He reports substantial improvements in three simple synthetic tasks.

Opitz and Shavlik (1996) take a similar approach, but they use a kind of genetic algorithm to search for a good population of neural network classifiers. In each iteration, they apply genetic operators to the current ensemble to generate new network topologies. These networks are then trained using an error function that combines the usual squared output prediction error with a multiplicative term that incorporates the diversity of the classifiers. After training the networks, they prune the population to retain the N best networks using a criterion that considers both accuracy and diversity. In a com-

parison with bagging, they found that their method gave excellent results in four real-world domains.

Abu-Mostafa (1990) and Caruana (1996) describe a technique for training a neural network on auxiliary tasks as well as the main task. The key idea is to add some output units to the network whose role is to predict the values of the auxiliary tasks and to include the prediction error for these auxiliary tasks in the error criterion that back propagation seeks to minimize. Because these auxiliary output units are connected to the same hidden units as the primary task output, the auxiliary outputs can influence the behavior of the network on the primary task. Parmanto et al. (1994) show that diverse classifiers can be learned by training on the same primary task but with many different auxiliary tasks. One good source of auxiliary tasks is to have the network attempt to predict one of its input features (in addition to the primary output task). They apply this method to medical diagnosis problems.

More recently, Munro and Parmanto (1997) developed an approach in which the value of the auxiliary output is determined dynamically by competition among the set of networks. Each network has a primary output y and a secondary (auxiliary) output z . During training, each network looks at the training example \mathbf{x} and computes its primary and secondary output predictions. The network whose secondary output prediction is highest is said to be the winner for this example. It is given a target value for z of 1; the remaining networks are given a target value of 0 for the secondary output. All networks are given the value y_i as the target for the primary output. The effect is to encourage different networks to become experts at predicting the secondary output z in different regions of the input space. Because the primary and secondary output share the hidden layer, the errors in the primary output become decorrelated. They show that this method substantially outperforms an ensemble of ordinary networks trained using different initial random weights when trained on a synthetic classification task.

In addition to these methods for training ensembles of neural networks, there are also methods that are specific to decision trees. Buntine (1990) developed an algorithm for learning *option trees*. These are decision trees where an internal node can contain several alternative splits (each producing its own sub-decision tree). To classify an example, each of these sub-decision trees is evaluated, and the resulting classifications are voted. Kohavi and Kunz (1997) describe an option tree algorithm and compare its performance to bagged C4.5 trees.

They show that option trees generally match the performance of bagging and produce a much more understandable result.

This completes my review of methods for generating ensembles using a single learning algorithm. Of course, one can always generate an ensemble by combining classifiers constructed by different learning algorithms. Learning algorithms based on very different principles will probably produce very diverse classifiers. However, often some of these classifiers perform much worse than others. Furthermore, there is no guarantee of diversity. Hence, when classifiers from different learning algorithms are combined, they should be checked (for example, by cross-validation) for accuracy and diversity, and some form of weighted combination should be used. This approach has been shown effective in some applications (for example, Zhang, Mesirov, and Waltz [1992]).

Methods for Combining Classifiers

Given that we have trained an ensemble of classifiers, how should we combine their individual classification decisions? Many methods have been explored. They can be subdivided into unweighted vote, weighted vote, and gating networks.

The simplest approach is to take an unweighted vote, as is done in bagging, ECOC, and many other methods. Although it might appear that more intelligent voting schemes should do better, the experience in the forecasting literature has been that simple, unweighted voting is robust (Clemen 1989). One refinement on simple majority vote is appropriate when each classifier h_ℓ can produce class-probability estimates rather than a simple classification decision. A *class-probability estimate* for data point \mathbf{x} is the probability that the true class is k : $P(f(\mathbf{x}) = k | h_\ell)$, for $k = 1, \dots, K$. We can combine the class probabilities of all the hypotheses so that the class probability of the ensemble is:

$$P(f(\mathbf{x}) = k) = \frac{1}{L} \sum_{\ell=1}^L P(f(\mathbf{x}) = k | h_\ell).$$

The predicted class of \mathbf{x} is then the class having the highest class probability.

Many different weighted voting methods have been developed for ensembles. For regression problems, Perrone and Cooper (1993) and Hashem (1993) apply least squares regression to find weights that maximize the accuracy of the ensemble on the training data. They show that the weight applied to h_ℓ should be inversely proportional to the variance of the estimates of h_ℓ . A difficulty with applying linear least squares is that the various hypotheses h_ℓ can be correlated highly. They describe methods for

choosing less correlated subsets of the ensemble and combining them by linear least squares.

For classification problems, weights are usually obtained by measuring the accuracy of each individual classifier h_ℓ on the training data (or a holdout data set) and constructing weights that are proportional to these accuracies. Ali and Pazzani (1996) describe a method that they call *likelihood combination*, in which they apply the naive Bayes algorithm (see The Naive Bayes Classifier) to learn weights for classifiers. In ADABOOST, the weight for classifier h_ℓ is computed from the accuracy of h_ℓ measured on the weighted training distribution that was used to learn h_ℓ . A Bayesian approach to weighted vote is to compute the posterior probability of each h_ℓ . This method requires the definition of a prior distribution $P(h_\ell)$ that is multiplied with the likelihood $P(S|h_\ell)$ to estimate the posterior probability of each h_ℓ . Ali and Pazzani experiment with this method. Earlier work on Bayesian voting of decision trees was performed by Buntine (1990).

The third approach to combining classifiers is to learn a gating network or a gating function that takes as input \mathbf{x} and produces as output the weights w_ℓ to be applied to compute the weighted vote of the classifiers h_ℓ . Jordan and Jacobs (1994) learn gating networks that have the form

$$w_\ell = e^{z_\ell} / \sum_u e^{z_u}$$

$$z_\ell = \mathbf{v}_\ell^T \mathbf{x}.$$

In other words, z_ℓ is the dot product of a parameter vector \mathbf{v}_ℓ with the input feature vector \mathbf{x} . The output weight w_ℓ is then the so-called *softmax* of the individual z_ℓ 's. As with any learning algorithm, there is a risk of overfitting the training data by learning the gating function in addition to learning each of the individual classifiers. (Below, I discuss the hierarchical mixture-of-experts method developed by Jordan and Jacobs that learns the w_ℓ and the gating network simultaneously.)

A fourth approach to combining classifiers, called *stacking*, works as follows: Suppose we have L different learning algorithms A_1, \dots, A_L and a set S of training examples $\{(x_1, y_1), \dots, (x_m, y_m)\}$. As usual, we apply each of these algorithms to our training data to produce hypotheses h_1, \dots, h_L . The goal of stacking is to learn a good combining classifier h^* such that the final classification will be computed by $h^*(h_1(\mathbf{x}), \dots, h_L(\mathbf{x}))$. Wolpert (1992) proposed the following scheme for learning h^* using a form of leave-one-out cross-validation.

Let

I already gave the basic intuition for why ensembles can improve performance: Uncorrelated errors made by the individual classifiers can be removed by voting. However, there is a deeper question lurking here: Why should it be possible to find ensembles of classifiers that make uncorrelated errors?

$$h^{(i)}$$

be a classifier constructed by algorithm A_i applied to all the training examples in S except example i . In other words, each algorithm is applied to the training data m times, leaving out one training example each time. We can then apply each classifier

$$h^{(i)}$$

to example x_i to obtain the predicted class

$$\hat{y}_i^l$$

This procedure gives us a new data set containing level-two examples whose features are the classes predicted by each of the L classifiers. Each example has the form

$$\langle (\hat{y}_i^1, \hat{y}_i^2, \dots, \hat{y}_i^L), y_i \rangle.$$

Now, we can apply some other learning algorithm to this level-two data to learn h^* . Breiman (1996b) applied this approach to combining different forms of linear regression with good results.

Why Ensembles Work

I already gave the basic intuition for why ensembles can improve performance: Uncorrelated errors made by the individual classifiers can be removed by voting. However, there is a deeper question lurking here: Why should it be possible to find ensembles of classifiers that make uncorrelated errors? There is another question as well: Why shouldn't we be able to find a single classifier that performs as well as the ensemble?

There are at least three reasons why good ensembles can be constructed and why it might be difficult or impossible to find a single classifier that performs as well as the ensemble. To understand these reasons, we must consider the nature of machine-learning algorithms. Machine-learning algorithms work by searching a space of possible hypotheses H for the most accurate hypothesis (that is, the hypothesis that best approximates the unknown function f). Two important aspects of the hypothesis space H are its size and the possibility it contains good approximations to f .

If the hypothesis space is large, then we need a large amount of training data to constrain the search for good approximations. Each training example rules out (or makes less plausible) all those hypotheses in H that misclassify it. In a two-class problem, ideally each training example can eliminate half of the hypotheses in H ; so, we require $O(\log|H|)$ examples to select a unique classifier from H .

The first cause of the need for ensembles is that the training data might not provide sufficient information for choosing a single best

classifier from H . Most of our learning algorithms consider very large hypothesis spaces; so, even after eliminating hypotheses that misclassify training examples, there are many hypotheses remaining. All these hypotheses appear equally accurate with respect to the available training data. We might have reasons for preferring some of these hypotheses over others (for example, preferring simpler hypotheses or hypotheses with higher prior probability), but nonetheless, there are typically many plausible hypotheses. From this collection of surviving hypotheses in H we can easily construct an ensemble of classifiers and combine them using the methods described earlier.

A second cause of the need for ensembles is that our learning algorithms might not be able to solve the difficult search problems that we pose. For example, the problem of finding the smallest decision tree that is consistent with a set of training examples is NP-hard (Hyafil and Rivest 1976). Hence, practical decision tree algorithms use search heuristics to guide a greedy search for small decision trees. Similarly, finding the weights for the smallest possible neural network consistent with the training examples is also NP-hard (Blum and Rivest 1988). Neural network algorithms therefore use local search methods (such as gradient descent) to find locally optimal weights for the network. A consequence of these imperfect search algorithms is that even if the combination of our training examples and our prior knowledge (for example, preferences for simple hypotheses, Bayesian priors) determines a unique best hypothesis, we might not be able to find it. Instead, we typically find a hypothesis that is somewhat more complex (or has somewhat lower posterior probability). If we run our search algorithms with a slightly different training sample or injected noise (or any of the other techniques described earlier), we find a different (suboptimal) hypothesis. Therefore, ensembles can be seen as a way of compensating for imperfect search algorithms.

A third cause of the need for ensembles is that our hypothesis space H might not contain the true function f . Instead, H might include several equally good approximations to f . By taking weighted combinations of these approximations, we might be able to represent classifiers that lie outside H . One way to understand this point is to visualize the decision boundaries constructed by learning algorithms. A *decision boundary* is a surface such that examples that lie on one side of the surface are assigned to a different class than examples that lie on the other side of the surface. The decision boundaries constructed by decision tree-learn-

ing algorithms are line segments (or, more generally, hyperplane segments) parallel to the coordinate axes. If the true boundary between two classes is a diagonal line, then decision tree algorithms must approximate the diagonal by a “staircase” of axis-parallel segments (figure 6). Different bootstrap training samples (or different weighted samples created by ADABOOST) shift the locations of the staircase approximation, and by voting among these different approximations, it is possible to construct better approximations to the diagonal decision boundary.

Interestingly, these improved staircase approximations are equivalent to complex decision trees. However, these trees are so large that were we to include them in our hypothesis space H the space would be far too large for the available training data. Hence, we can see that ensembles provide a way of overcoming representational inadequacies in our hypothesis space.

Open Problems Concerning Ensembles

Ensembles are well established as a method for obtaining highly accurate classifiers by combining less accurate ones. There are still many questions, however, about the best way to construct ensembles as well as issues about how best to understand the decisions made by ensembles.

Faced with a new learning problem, what is the best approach to constructing and applying an ensemble of classifiers? In principle, there can be no single best ensemble method, just as there can be no single best learning algorithm. However, some methods might be uniformly better than others, and some methods might be better than others in certain situations.

Experimental studies have shown that ADABOOST is one of the best methods for constructing ensembles of decision trees. Schapire (1997) compares ADABOOST.M2 and ADABOOST.OC to bagging and error-correcting output coding and shows that the ADABOOST methods are generally superior. However, Quinlan (1996) has shown that in domains with noisy training data, ADABOOST.M1 can perform badly—it places high weight on incorrectly labeled training examples and, consequently, constructs bad classifiers. Kong and Dietterich (1995) showed that combining bagging with error-correcting output coding improved the performance of both methods, which suggests that combinations of other ensemble methods should be explored as well. Dietterich and Kong also showed that error-correcting output coding does not work well with highly local algorithms (such as




Figure 6. Decision Boundaries.

a. This figure shows the true diagonal decision boundary and three staircase approximations to it (of the kind that are created by decision tree algorithms). b. This figure shows the voted decision boundary, which is a much better approximation to the diagonal boundary.

nearest-neighbor methods).

There have been few systematic studies of methods for constructing ensembles of neural networks, rule-learning systems, and other types of classifier. Much work remains in this area.

Although ensembles provide accurate classifiers, some problems may limit their practical application. One problem is that ensembles can require large amounts of memory to store and large amounts of computation to apply. For example, earlier I mentioned that an ensemble of 200 decision trees attains perfect performance on a letter-recognition benchmark task. However, these 200 decision trees require 59 megabytes of storage, which makes them impractical for most present-day computers. An important line of research, therefore, is to find ways of converting these ensembles into less redundant representations, perhaps through the deletion of highly correlated members of the ensemble or by representational transformations.

A second difficulty with ensemble classifiers is that an ensemble provides little insight into how it makes its decisions. A single decision tree can often be interpreted by human users, but an ensemble of 200 voted decision trees is much more difficult to understand. Can methods be found for obtaining explanations (at least locally) from ensembles? One example of work on this question is Craven’s TREPAN algorithm (Craven and Shavlik 1996).

Scaling Up Machine-Learning Algorithms

A second major research area has explored techniques for scaling up learning algorithms so that they can apply to problems with millions of training examples, thousands of features, and hundreds of classes. Large machine-

Large machine-learning problems are beginning to arise in database-mining applications, where there can be millions of transactions every day and where it is desirable to have machine-learning algorithms that can analyze such large data sets in just a few hours of computer time.

learning problems are beginning to arise in database-mining applications, where there can be millions of transactions every day and where it is desirable to have machine-learning algorithms that can analyze such large data sets in just a few hours of computer time. Another area where large learning problems arise is in information retrieval from full-text databases and the World Wide Web. In information retrieval, each word in a document can be treated as an input feature; so, each training example can be described by thousands of features. Finally, applications in speech recognition, object recognition, and character recognition for Chinese and Japanese present situations in which hundreds or thousands of classes must be discriminated.

Learning with Large Training Sets

Decision tree algorithms have been extended to handle large data sets in three different ways. One approach is based on intelligently sampling subsets of the training data as the tree is grown. To describe how this process works, I must review how decision tree algorithms operate.

Decision trees are constructed by starting with the entire training set and an empty tree. A test is chosen for the root of the tree, and the training data are then partitioned into disjoint subsets depending on the outcome of the test. The algorithm is then applied recursively to each of these disjoint subsets. The algorithm terminates when all (or most) of the training examples within a subset of the data belong to the same class. At this point, a leaf node is created and labeled with the class.

The process of choosing a test for the root of the tree (or for the root of each subtree) involves analyzing the training data and choosing the one feature that is the best predictor of the output class. In a large and redundant data set, it might be possible to make this choice based on only a sample of the data. Musick, Catlett, and Russell (1993) presented an algorithm that dynamically chooses the sample based on how difficult the decision is at each node. The algorithm typically behaves by using a small sample near the root of the tree and then progressively enlarging the sample as the tree grows. This technique can reduce the time required to grow the tree without reducing the accuracy of the tree at all.

A second approach is based on developing clever data structures that avoid the need to store all training data in random-access memory. The hardest step for most decision tree algorithms is to find tests for real-valued features. These tests usually take the form $x_j \leq \theta$

for some threshold value θ . The standard approach is to sort the training examples at the current node according to the values of feature x_j and then make a sequential pass over the sorted examples to choose the threshold θ . In standard depth-first algorithms, the data must be sorted by each candidate feature x_j at each node of the tree, which can become expensive.

Shafer, Agrawal, and Mehta (1996) describe their SPRINT method in which the training data are broken up into a separate disk file for each attribute (and sorted by attribute value). For feature j , the disk file contains records of the form $\langle i, x_{ij}, y_i \rangle$, where i is the index of the training example, x_{ij} is the value of feature j for training example i , and y_i is the class of example i . To choose a splitting threshold θ for feature j , it is a simple matter to make a serial scan of this disk file and construct class histograms for each potential splitting threshold. Once a value is chosen, the disk file is partitioned (logically) into two files, one containing examples with values less than or equal to θ and the other containing examples with values greater than θ . As these files are written to disk, a hash table is built (in the main memory) in which the index for each training example i is associated with the left or right child nodes of the newly created split. Then, each of the disk files for the other attributes $x_l, l \neq j$ is read and split into a file for the left child and a file for the right child. The index of each example is looked up in the hash table to determine the child to which it belongs. SPRINT can be parallelized easily, and it has been applied to data sets containing 2.5 million examples. Mehta, Agrawal, and Rissanen (1996) have developed a closely related algorithm, SLIQ, that makes more use of main memory but also scales to millions of training examples and is slightly faster than SPRINT. Both SPRINT and SLIQ scale approximately linearly in the number of training examples and the number of features, aside from the cost of performing the initial sorting of the data (which need only be done once).

A third approach to large data sets is to take advantage of ensembles of decision trees (Chan and Stolfo 1995). The training data can randomly be partitioned into N disjoint subsets. A separate decision tree can be grown from each subset in parallel. The trees can then vote to make classification decisions. Although the accuracy of each of the individual decision trees is less than the accuracy of a single tree grown with all the data, the accuracy of the ensemble is often better than the accuracy of a single tree. Hence, with N parallel processors, we can achieve a speedup of N in the time required to construct the decision trees.

A fourth approach to the problem of choosing splits for real-valued input features is to discretize the values of such features. For example, one feature of a training example might be employee income measured in dollars. This could take on tens of thousands of distinct values, and the running time of most decision-tree learning algorithms is linear in the number of distinct values of each feature. This problem can be solved by grouping income into a small number of ranges (for example, \$0–\$10,000; \$10,000–\$25,000; \$25,000–\$50,000; \$50,000–\$100,000; and greater than \$100,000). If the ranges are chosen well, the resulting decision trees will still be accurate. Several simple and fast algorithms have been developed for choosing good discretization points (Kohavi and Sahami 1996; Fayyad and Irani 1993; Catlett 1991).

A very effective rule-learning algorithm, called RIPPER, has been developed by William Cohen (1995) based on an earlier algorithm, IREP, developed by Furnkranz and Widmer (1994). RIPPER constructs rules of the form $test_1 \wedge test_2 \wedge \dots \wedge test_t \Rightarrow c$, where each test has the form $x_j = \theta$ (for discrete features) or $x_j \leq \theta$ or $x_j > \theta$ (for real-valued features). A rule is said to cover a training example if the example satisfies all the tests on the left-hand side of the rule. For a training set of size m , the run time of RIPPER scales as $O(m(\log m)^2)$. This is a major improvement over the rule-learning program C4.5RULES (Quinlan 1993), which scales as $O(m^3)$.

Figure 7 shows pseudocode for RIPPER. RIPPER works by building an initial set of rules and optimizing the set of rules k times, where k is a parameter (typically set to 2). I describe RIPPER for the case where there are only two classes, 1 and 2. Examples of class 1 will be referred to as the positive examples, and examples of class 2 will be referred to as the negative examples. RIPPER can easily be extended to handle larger numbers of classes.

To build a set of rules, RIPPER constructs one rule at a time. Before learning each rule, it divides the training data into a growing set (containing two-thirds of the data) and a pruning set (containing the remaining one-third). It then iteratively adds tests to the rule until the rule covers no negative examples. Tests are selected using an information-gain heuristic developed for Quinlan's (1990) FOIL system. Once a rule is grown, it is immediately pruned by deleting tests in reverse order, $test_t, test_{t-1}, \dots, test_1$, to find the pruned rule that maximizes the quantity $(p - n)/(p + n)$, where p is the number of positive pruning examples covered by the pruned rule, and n is the number of negative

```

procedure BUILDRULESET( $P, N$ )
 $P$  = positive examples
 $N$  = negative examples
 $RuleSet = \{ \}$ 
 $DL = DescriptionLength(RuleSet, P, N)$ 
while  $P \neq \{ \}$ 
  // Grow and prune a new rule
  split ( $P, N$ ) into ( $GrowPos, GrowNeg$ ) and ( $PrunePos, PruneNeg$ )
   $Rule := GrowRule(GrowPos, GrowNeg)$ 
   $Rule := PruneRule(Rule, PrunePos, PruneNeg)$ 
  add  $Rule$  to  $RuleSet$ 
  if  $DescriptionLength(RuleSet, P, N) > DL + 64$  then
    // Prune the whole rule set and exit
    for each rule  $R$  in  $RuleSet$  (considered in reverse order)
      if  $DescriptionLength(RuleSet - \{R\}, P, N) < DL$  then
        delete  $R$  from  $RuleSet$ 
         $DL := DescriptionLength(RuleSet, P, N)$ 
      end if
    end for
  return ( $RuleSet$ )
end if
 $DL := DescriptionLength(RuleSet, P, N)$ 
delete from  $P$  and  $N$  all examples covered by  $Rule$ 
end while
end BUILDRULESET

procedure OPTIMIZERULESET( $RuleSet, P, N$ )
for each rule  $R$  in  $RuleSet$ 
  delete  $R$  from  $RuleSet$ 
   $UPos :=$  examples in  $P$  not covered by  $RuleSet$ 
   $UNeg :=$  examples in  $N$  not covered by  $RuleSet$ 
  split ( $UPos, UNeg$ ) into ( $GrowPos, GrowNeg$ ) and ( $PrunePos, PruneNeg$ )
   $RepRule := GrowRule(GrowPos, GrowNeg)$ 
   $RepRule := PruneRule(RepRule, PrunePos, PruneNeg)$ 
   $RevRule := GrowRule(GrowPos, GrowNeg, R)$ 
   $RevRule := PruneRule(RevRule, PrunePos, PruneNeg)$ 
  choose better of  $RepRule$  and  $RevRule$  and add to  $RuleSet$ 
end for
end OPTIMIZERULESET

procedure RIPPER( $P, N, k$ )
 $RuleSet := BUILDRULESET(P, N)$ 
repeat  $k$  times  $RuleSet := OPTIMIZERULESET(RuleSet, P, N)$ 
return ( $RuleSet$ )
end RIPPER

```

Figure 7. The RIPPER Algorithm (Cohen 1995).

pruning examples covered by the pruned rule.

Once the rule has been grown and pruned, RIPPER adds it to the rule set and discards all training examples that are covered by this new rule. It uses a description-length criterion to decide when to stop adding rules. The *description length* of a set of rules is the number of bits needed to represent the rules plus the number of bits needed to identify the training examples that are exceptions to the rules. Minimum-description-length criteria of this kind have been applied successfully to rule- and tree-learning algorithms (Quinlan and Rivest 1989). RIPPER stops adding rules when the description length of the rule set is more than 64 bits larger than the best description length observed to this point. It then considers the rules in reverse order and deletes any rule that will reduce the total description length of the rule set.

To optimize a set of rules, RIPPER considers deleting each rule in turn and regrowing and

repruning it. Two candidate replacement rules are grown and pruned. The first candidate is grown starting with an empty rule, whereas the second candidate is grown starting with the current rule. The better of the two candidates is selected (using a description-length heuristic) and added to the rule set.

Cohen compared RIPPER to C4.5RULES on 37 data sets and found that RIPPER matched or beat C4.5RULES in 22 of the 37 problems. The rule sets that it finds are always smaller than those constructed by C4.5RULES. An implementation of RIPPER is available for research use from Cohen (www.research.att.com/~wcohen/ripperd.html).

Learning with Many Features

In many learning problems, there are hundreds or thousands of potential features describing each input object \mathbf{x} . Popular learning algorithms such as C4.5 and back propagation do not scale well when there are many features. Indeed, from a statistical point of view, examples with many irrelevant, but noisy, input features provide little information. It is easy for learning algorithms to be confused by the noisy features and construct poor classifiers. Hence, in practical applications, it is wise to carefully choose which features to provide to the learning algorithm.

Research in machine learning has sought to automate the selection and weighting of features, and many different algorithms have been developed for this purpose. An excellent review has been written by Wettschereck, Aha, and Mohri (1997). A comprehensive review of the statistical literature on feature selection can be found in Miller (1990). I discuss a few of the most significant methods here.

Three main approaches have been pursued. The first approach is to perform some initial analysis of the training data and select a subset of the features to feed to the learning algorithm. A second approach is to try different subsets of the features on the learning algorithm, estimate the performance of the algorithm with these features, and keep the subsets that perform best. The third approach is to integrate the selection and weighting of features directly into the learning algorithm. I discuss two examples of each approach.

Selecting and Weighting Features by Preprocessing A simple preprocessing technique is to compute the mutual information (also called the information gain) between each input feature and the class. The *mutual information* between two random variables is the average reduction in uncertainty about the second variable given a value of the

first. For discrete feature j , the mutual information weight w_j can be computed as

$$w_j = \sum_v \sum_c P(y = c, x_j = v) \cdot \log \frac{P(y = c, x_j = v)}{P(y = c)P(x_j = v)},$$

where $P(y = c)$ is the proportion of training examples in class c , and $P(x_j = v)$ is the probability that feature j takes on value v . For real-valued features, the sums become integrals that must be approximated. A good approximation is to apply a discretization algorithm, such as the one advocated by Fayyad and Irani (1993); convert the real-valued feature into a discrete-valued feature; and then apply the previous formula. Wettschereck and Dietterich (1995) have obtained good results with nearest-neighbor algorithms using mutual information weighting.

A problem with mutual information weighting is that it treats each feature independently. For features whose predictive power is only apparent in combination with other features, mutual information assigns a weight of zero. For example, a difficult class of learning problems involves learning parity functions with random irrelevant features. A parity function over n binary features is equal to 1 if and only if an odd number of the features are equal to 1. Suppose we define a learning problem in which there are 4 relevant features and 10 irrelevant (random) binary features, and the class is the 4-parity of the 4 relevant features. The mutual information weights of all features will be approximately zero using the previous formula.

An algorithm that overcomes this problem (and is one of the most successful preprocessing algorithms to date) is the RELIEF-F algorithm (Kononenko 1994), which is an extension of an earlier algorithm called RELIEF (Kira and Rendell 1992). The basic idea of these algorithms is to draw examples at random, compute their nearest neighbors, and adjust a set of feature weights to give more weight to features that discriminate the example from neighbors of different classes. Specifically, let \mathbf{x} be a randomly chosen training example, and let \mathbf{x}' and \mathbf{x}'' be the two training examples nearest to \mathbf{x} (in Euclidean distance) in the same class and in a different class, respectively. The goal of RELIEF is to set the weight w_j on input feature j to be

$$w_j = P(x_j \neq x_j') - P(x_j \neq x_j'').$$

In other words, the weight w_j should be maximized when \mathbf{x}_j has a high probability of taking on a different value for feature j , and \mathbf{x}'_j has a low probability of taking on a different value for feature j . RELIEF-F computes a more reliable estimate of this probability difference by computing the B nearest neighbors of \mathbf{x} in each class.

Figure 8 describes the RELIEF-F algorithm. In this algorithm, $\delta(u, v)$ for two feature values u and v is defined as follows:

$$\delta(u, v) = \begin{cases} |u - v| & \text{for real-valued features} \\ 0 & \text{if } u = v \\ 1 & \text{if } u \neq v \text{ for discrete features.} \end{cases}$$

Kononenko, Simec, and Robnik-Sikonja (1997) have shown that RELIEF-F is effective at detecting relevant features, even when these features are highly dependent on other features.

For the 4-parity problem mentioned previously (with 10 irrelevant random features), RELIEF-F can correctly separate the 4 relevant features from the 10 irrelevant ones given 400 training examples. Kononenko computes the 10 nearest neighbors in each class (that is, $B = 10$). In his experiments, he sets the number of sample points L to be equal to the number of training examples, but in large data sets, good results can be obtained from much smaller samples.

Kononenko et al. have also experimented with integrating RELIEF-F into a decision tree-learning algorithm called ASSISTANT-R. They show that ASSISTANT-R is able to perform much better than the original ASSISTANT program (which uses mutual information to choose features) in domains with highly dependent features but gives essentially the same performance in domains with independent features.

Selecting and Weighting Features by Testing with the Learning Algorithm

John, Kohavi, and Pfleger (1994) describe a computationally expensive method that they call the *wrapper method* for selecting input features. The idea is to generate sets of features, run the learning algorithm using only these features, and evaluate the resulting classifiers using 10-fold cross-validation (or a single hold-out set). In 10-fold cross-validation, the training data are subdivided randomly into 10 disjoint equal-sized sets. The learning algorithm is applied 10 times, each time on a training set containing all but one of these subsets. The resulting classifier is tested on the one-tenth of the data that was held out. The performance of the 10 classifiers (on their 10 respective hold-out sets) is averaged to provide an estimate of the overall performance of the learning algorithm when trained with the given features.

Kohavi and John explored step-wise selection algorithms that start with a set of features (for example, the empty set) and considered adding or deleting a single feature. The possible changes to the feature set are evaluated (using 10-fold cross-validation), and the best change is made. Then, a new set of changes is

```

procedure RELIEF-F( $L, B$ )
 $L$  = the number of random examples to draw
 $B$  = the number of near neighbors to compute
for all features  $j$ :  $w_j := 0.0$ 
 $p_c$  := the fraction of the training examples belonging to class  $c$ 
for  $l := 1$  to  $L$  do
  randomly select an instance  $(\mathbf{x}_l, y_l)$ 
  let  $H_{it}$  be the set of  $B$  examples  $(\mathbf{x}_i, y_i)$  nearest to  $\mathbf{x}_l$  such that  $y_i = y_l$ .
  for each class  $c \neq y_l$ 
    let  $M_c$  be a set of  $B$  examples  $(\mathbf{x}_i, y_i)$  nearest to  $\mathbf{x}_l$  such that  $y_i = c$ .
  end for
  for each feature  $j$ 
 $w_j := w_j - \frac{1}{LB} \sum_{(\mathbf{x}_i, y_i) \in H_{it}} \delta(x_{lj}, x_{ij}) + \sum_{c \neq y_l} \frac{p_c}{(1-p_c)LB} \sum_{(\mathbf{x}_i, y_i) \in M_c} \delta(x_{lj}, x_{ij})$ 
  end for  $j$ 
end for  $l$ 
return  $w_j \forall j$ .
end RELIEF-F

```

Figure 8. The RELIEF-F Algorithm.

considered. This method is only practical for data sets with relatively small numbers of features and fast learning algorithms, but it gave excellent results on the University of California at Irvine benchmarks.

Moore and Lee (1994) describe a much more efficient approach to feature selection that combines leave-one-out cross-validation (LOOCV) with the nearest-neighbor algorithm. In *leave-one-out cross-validation*, each training example is temporarily deleted from the training data, and the nearest-neighbor learning algorithm is applied to predict the class of the example. The total number of classification errors is the leave-one-out cross-validated estimate of the error rate of the learning algorithm. Moore and Lee use the LOOCV error to compare different sets of features with the goal of finding the set of relevant features that minimizes the LOOCV error.

They combine two clever ideas to achieve this. The first idea is called *racine*: Suppose we are considering two different sets of relevant features, A and B . We repeatedly choose a training example at random, temporarily delete it from the training set, and apply the nearest-neighbor rule to classify it using the features in set A and the features in set B . We count the number of classification errors corresponding to each set. As we process more and more training examples in this leave-one-out fashion, the error rate for feature set A can become so much larger than the error rate for B that we can conclude with high confidence that B is the better feature set and terminate the race. In their paper, Moore and Lee (1994) apply Bayesian statistics to make this termination decision.

The second idea is based on schemas. We can represent each set of relevant features by a

bit vector where a 1 in position j means that feature j is relevant, and a 0 means it is irrelevant. A schema is a vector containing 0s, 1s, and \otimes s. A \otimes in position j means that this feature should be selected randomly to be relevant 50 percent of the time. Moore and Lee race pairs of schemas against one another as follows: A training example is randomly selected and temporarily deleted from the training set. The nearest-neighbor algorithm is applied to classify it using each of the two schemas being raced. To classify an example using a schema, features indicated by a 0 are ignored, features indicated by a 1 are selected, and features indicated by a \otimes are selected with probability 0.5. Suppose, for illustration, that we have five features. Moore and Lee begin by conducting five simultaneous pairwise races:

1 $\otimes\otimes\otimes\otimes$ races against 0 $\otimes\otimes\otimes\otimes$
 \otimes 1 $\otimes\otimes\otimes$ races against \otimes 0 $\otimes\otimes\otimes$
 $\otimes\otimes$ 1 $\otimes\otimes$ races against $\otimes\otimes$ 0 $\otimes\otimes$
 $\otimes\otimes\otimes$ 1 \otimes races against $\otimes\otimes\otimes$ 0 \otimes
 $\otimes\otimes\otimes\otimes$ 1 races against $\otimes\otimes\otimes\otimes$ 0

All the races are terminated as soon as one of the schemas is found to be better than its opponent. In the next iteration, all single-bit refinements of the winning schema are raced against one another. For example, suppose schema 1 was the winner of the first race. Then, the next iteration involves the following four pairwise races:

11 $\otimes\otimes\otimes$ races against 01 $\otimes\otimes\otimes$
 \otimes 11 $\otimes\otimes$ races against \otimes 10 $\otimes\otimes$
 \otimes 1 \otimes 1 \otimes races against \otimes 1 \otimes 0 \otimes
 \otimes 1 $\otimes\otimes$ 1 races against \otimes 1 $\otimes\otimes$ 0

This process continues until all the \otimes s are removed from the winning schema. Moore and Lee found that this method never misses important relevant features, even when the features are highly dependent. However, in some rare cases, the races can take a long time to conclude. The problem appears to be that the algorithm can be slow to replace a \otimes with a 0. In contrast, the algorithm is quick to replace a \otimes with a 1. Moore and Lee, therefore, investigated an algorithm, called SCHEMATA+, that terminates each race after 2000 evaluations (in favor of the 0, using the race statistics to choose the feature least likely to be relevant). The median number of training examples that must be evaluated by SCHEMATA+ is only 13 percent of the number of evaluations required by greedy forward selection and only 11 percent of the number of evaluations required by greedy backward elimination, but it achieves the same levels of accuracy. An important direction for future research is to compare the accuracy and speed of SCHEMATA+ and RELIEF-F

to see which works better with various learning algorithms.

Integrating Feature Weighting into the Learning Algorithm

I now discuss two methods that integrate feature selection directly into the learning algorithm. Both of them have been shown to work well experimentally, and the second method, called WINNOW, works extremely well in problems with thousands of potentially relevant input features.

The first algorithm is called the variable-kernel similarity metric, or vsm method (Lowe 1995). vsm is a form of Gaussian radial basis function method. To classify a new data point \mathbf{x}_r , it defines a multivariate Gaussian probability distribution ϕ centered on \mathbf{x}_r with standard deviation σ . Each example (\mathbf{x}_i, y_i) in the training data set “votes” for class y_i with an amount $\phi(\|\mathbf{x}_i - \mathbf{x}_r\|; \sigma)$. The class with the highest vote is assigned to be the class $y_r = f(\mathbf{x}_r)$ of the data point \mathbf{x}_r .

The key to the effectiveness of vsm is that it learns a weighted distance metric for measuring the distance between the new data point \mathbf{x} and each training point \mathbf{x}_r . vsm also adjusts the size σ of the Gaussian distribution depending on the local density of training examples in the neighborhood of \mathbf{x}_r .

In detail, vsm is controlled by a set of learned feature weights w_1, \dots, w_n , a kernel radius parameter r , and a number of neighbors R . To classify a new data point \mathbf{x}_r , vsm first computes the weighted distances to the R nearest neighbors (\mathbf{x}_i, y_i) :

$$d_i = \sqrt{\sum_{j=1}^n w_j^2 (x_{ij} - x_{rj})^2}$$

It then computes a kernel width σ from the average distance to the $R/2$ nearest neighbors:

$$\sigma = \frac{2r}{R} \sum_{i=1}^{R/2} d_i.$$

Finally, it computes the probability that the next point \mathbf{x} belongs to each class c (the quantity v_i is the vote from training example i):

$$v_i = \exp(-d_i^2 / 2\sigma)$$

$$P[f(\mathbf{x}_r) = c] = \frac{\sum_{i=1}^R v_i [y_i = c]}{\sum_{i=1}^R v_i}.$$

vsm then guesses the class with the highest probability.

How does vsm learn the values of the feature weights and the kernel radius r ? It does so by performing gradient-descent search to minimize the LOOCV accuracy of the vsm classifier. Lowe (1995) shows how to compute the gradient of the LOOCV error with respect to each of the weights w_j and the parameter r . Starting

with initial values for these parameters, VSM computes the R nearest neighbors of each training example (\mathbf{x}, y) . It then computes the gradient and performs a search in the direction of the gradient to minimize LOOCV error (and keeps this set of nearest neighbors fixed). The search along the direction of the gradient is called a *line search*, and there are several efficient algorithms available (Press et al. 1992). Even though the weights are changing during the line search, the set of nearest neighbors (and the gradient) is not recomputed. Once the error is minimized along this direction of the gradient, the R nearest neighbors are recomputed, a new gradient is computed, and a new line search is performed. Lowe applied the conjugate gradient algorithm to select each new search direction, and he reports that generally only 5 to 30 line searches were required to minimize the LOOCV error. The resulting classifiers gave excellent results on two challenging benchmark tasks.

The last feature weighting algorithm I discuss is the WINNOW algorithm developed by Littlestone (1988). WINNOW is a linear threshold algorithm for two-class problems with binary (that is, 0/1 valued) input features. It classifies a new example \mathbf{x} into class 2 if

$$\sum_i w_i x_i > \theta$$

and into class 1 otherwise. WINNOW is an online algorithm; it accepts examples one at a time and updates the weights w_j as necessary. Pseudocode for the algorithm is shown in figure 9.

WINNOW initializes its weights w_j to 1. It then accepts a new example (\mathbf{x}, y) and applies the threshold rule to compute the predicted class y' . If the predicted class is correct ($y' = y$), WINNOW does nothing. However, if the predicted class is wrong, WINNOW updates its weights as follows. If $y' = 0$ and $y = 1$, then the weights are too low; so, for each feature such that $x_j = 1$, $w_j := w_j \cdot \alpha$, where α is a number greater than 1 called the *promotion parameter*. If $y' = 1$ and $y = 0$, then the weights were too high; so, for each feature $x_j = 1$, it decreases the corresponding weight by setting $w_j := w_j \cdot \beta$, where β is a number less than 1 called the *demotion parameter*.

The fact that WINNOW leaves the weights unchanged when the predicted class is correct is somewhat puzzling to many people. It turns out to be critical for its successful behavior, both theoretically and experimentally. One explanation is that like ADABOOST, this strategy focuses WINNOW's attention on its mistakes. Another explanation is that it helps prevent overfitting.

Littlestone's theoretical analysis of WINNOW introduced the worst-case mistake bound

```

procedure WINNOW( $\alpha, \beta, \theta$ )
 $\alpha > 1$  is the promotion parameter
 $\beta < 1$  is the demotion parameter
 $\theta$  is the threshold
initialize  $w_j := 1$  for all  $j$ 
for each training example  $(\mathbf{x}, y)$ 
   $z := \sum_j w_j \cdot x_j$ 
   $y' := \begin{cases} 0 & \text{if } z < \theta \\ 1 & \text{if } z \geq \theta \end{cases}$ 
  if  $y' = 0$  and  $y = 1$  then
     $w_j := w_j \cdot \alpha$  for each  $j$  such that  $x_j = 1$ 
  end if
  else if  $y' = 1$  and  $y = 0$  then
     $w_j := w_j \cdot \beta$  for each  $j$  such that  $x_j = 1$ 
  end if
end for
end WINNOW

```

Figure 9. The WINNOW Algorithm.

method. The idea is to assume that the true function $y = f(\mathbf{x})$ belongs to some set of classifiers H and derive a bound on the maximum number of mistakes that WINNOW makes when an adversary is allowed to choose f and the order in which the training examples are presented to WINNOW. Littlestone proves the following result:

Theorem 1: Let f be a disjunction of r out of its n input features. Then, WINNOW will learn f and make no more than $2 + 3r(1 + \log_2 n)$ mistakes (for $\alpha = 2$, $\beta = 1/2$, and $\theta = n$).

This theorem shows that the convergence time of WINNOW is linear in the number of relevant features r and only logarithmic in the total number of features n . Similar results hold for other values of the α , β , and θ parameters, which permits WINNOW to learn functions where u out of v of the features must be 1 and many other interesting classes of Boolean functions.

WINNOW is an example of an *exponential update algorithm*. The weights of the relevant features grow exponentially, but the weights of the irrelevant features shrink exponentially. General results in computational learning theory have developed similar exponential update algorithms for many applications. A common property of these algorithms is that they excel when the number of relevant features is small compared to the total number of features.

WINNOW has been applied to several experimental learning problems. Blum (1997) describes an application to a calendar scheduling task. In this task, a calendar system is given a description of a proposed meeting (including the list of invitees and their properties). It must then predict the start time, day of week, location, and duration of the meeting. There were




Figure 10. Comparison of the Percentage of Correct Classifications for a Modified Bayesian Method and WINNOW for 21 Sets of Frequently Confused Words.

Trained on 80 percent of the Brown corpus and tested on the remaining 20 percent. Points lying above the line $y = x$ correspond to cases where WINNOW is more accurate.

34 input features available, some with many possible values. Blum defined Boolean features corresponding to all possible values of all pairs of the input features. For example, given two input features—event-type and position-of-attendees—he would define a separate Boolean feature for each legal combination, such as event-type = meeting and position-of-attendees = grad-student. There were a total of 59,731 Boolean input features. He then applied WINNOW with $\alpha = 3/2$ and $\beta = 1/2$. He modified WINNOW to prune (set to 0) weights that become very small (less than 0.00001).

WINNOW found that 561 of the Boolean features were actually useful for prediction. Its accuracy was better than that of the best previous classifier for this task (which used a greedy forward-selection algorithm to select relevant features for a decision tree-learning algorithm).

Golding and Roth (1996) describe an application of WINNOW to context-sensitive spelling correction. This is the task of identifying spelling errors where one legal word is substituted for another, such as “It’s not to late,” where *to* is substituted for *too*. The Random House dictionary (Flexner 1983) lists many sets of commonly confused words, and Golding and Roth developed a separate WINNOW classifier for each of the listed sets (for exam-

ple, {to, too, two}). WINNOW’s task is to decide whether each occurrence of these words is correct or incorrect based on its context.

Golding and Roth use two kinds of Boolean input feature. The first kind are context words. A *context word* is a feature that is true if a particular word (for example, “cloudy”) appears within 10 words before or after the target word. The second kind are *collocation features*. These test for a string of two words or part-of-speech tags immediately adjacent to the target word. For example, the sequence *target to VERB* is a collocation feature that checks whether the target word is immediately followed by the word *to* and then a word that can potentially be a verb (according to a dictionary lookup). Based on the one-million-word Brown corpus (Kucera and Francis 1967), Golding and Roth defined more than 10,000 potentially relevant features.

Golding and Roth applied WINNOW (with $\alpha = 3/2$, β varying between 0.5 and 0.9, and $\theta = 1$). They compared its accuracy to the best previous method, which is a modified naive Bayesian algorithm. Figure 10 shows the results for 21 sets of frequently confused words.

An important advantage of WINNOW, in addition to its speed and ability to operate online, is that it can adapt rapidly to changes in the target function. This advantage was shown to be important in the calendar scheduling problem, where the scheduling of different types of meeting can shift when semesters change (or during summer break). Blum’s experiments showed that WINNOW was able to respond quickly to such changes. By comparison, to enable the decision tree algorithm to respond to changes, it was necessary to decide which old training examples could be deleted. This task is difficult to do because although some kinds of meeting might change with the change in semesters, other meetings might stay the same. A decision to keep, for example, only the most recent 180 training examples means that examples of rare meetings whose scheduling does not change will be lost, which hurts the performance of the decision tree approach. In contrast, because WINNOW only revises the weights on features when they have the value 1, features describing such rare meetings are likely to retain their weights even as the weights for other features are being modified rapidly.

Summary: Scaling Up Learning Algorithms

This concludes my review of methods for scaling up learning algorithms to apply to very large problems. With the techniques described here, problems having one million training examples can be solved in reasonable amounts of

computer time. However, it is not clear whether the current stock of ideas will permit the solution of problems with billions of training examples. An important open problem is to gather more practical experience with very large problems, so that we can understand their properties and determine where these algorithms fail.

A recurring theme is the use of subsamples of the training data to make critical intermediate decisions (such as the choice of relevant features). Another theme is the development of efficient online algorithms, such as WINNOWER. These are anytime algorithms that can produce a useful answer regardless of how long they are permitted to run. The longer they run, the better the result they produce.

An important open topic is the problem of handling thousands of output classes. The section Ensembles of Classifiers has already described the two methods that are most appropriate in this case: error-correcting output coding and ADABOOST.OC. Both of these methods should scale well with the number of classes. Error-correcting output coding has been tested on problems with as many as 126 classes, but tests on very large problems with thousands of classes have not yet been performed.

Reinforcement Learning

The previous two sections discussed problems in supervised learning from examples. This section addresses problems of sequential decision making and control that come under the heading of reinforcement learning.

Work in reinforcement learning dates back to the earliest days of AI when Arthur Samuel (1959) developed his famous checkers program. More recently, there have been several important advances in the practice and theory of reinforcement learning. Perhaps the most famous work is Gerry Tesauro's (1992) TD-GAMMON program, which has learned to play backgammon better than any other computer program and almost as well as the best human players. Two other interesting applications are the work of Zhang and Dietterich (1995) on job-shop scheduling and Crites and Barto (1995) on real-time scheduling of passenger elevators.

Kaelbling, Littman, and Moore (1996) published an excellent survey of reinforcement learning, and Mahadevan and Kaelbling (1996) report on a recent National Science Foundation-sponsored workshop on the subject. Two new books (Barto and Sutton 1997; Bertsekas and Tsitskilis 1996) describe the newly developed reinforcement learning algorithms and the theory behind them. I summarize these developments here.

An Introduction to Dynamic Programming

The most important insight of the past five years is that reinforcement learning is best analyzed as a form of online, approximate dynamic programming (Barto, Bradtke, and Singh 1995). I introduce this insight using the following notation: Consider a robot interacting with an external environment. At each time t , the environment is in some state s_t , and the robot has available some set of actions A . The robot executes an action a_t , which causes the environment to move to a new state s_{t+1} . A convenient way of specifying the desired behavior of the robot is to define an immediate reward function $R(s_t, a, s_{t+1})$ that specifies a real-valued reward for this transition from s_t to s_{t+1} . For example, we can assign a positive reward to actions that reach a desired goal location, and we can assign a negative reward to undesirable actions such as colliding with walls, people, or other robots. The immediate reward in all other states could be defined as 0.

Our long-term goal for the robot can then be defined as some function of the immediate rewards it receives. A commonly used criterion is the *cumulative discounted reward*,

$$\sum_{t=0}^{\infty} \gamma^t R(s_t, a, s_{t+1}),$$

where $0 \leq \gamma < 1$ is a *discount factor* that controls the relative importance of short-term and long-term rewards.

A procedure or rule for choosing each action a given state s is called the *policy* of the robot, and it can be formalized as a function $a = \pi(s)$. The goal of reinforcement learning algorithms is to compute the *optimal policy*, denoted π^* , which maximizes the cumulative discounted reward.

Researchers in dynamic programming (for example, Bellman [1957]) found it convenient to define a real-valued function $f^\pi(s)$ called the *value function* of policy π . The value function $f^\pi(s)$ gives the expected cumulative discounted reward that will be received by starting in state s and executing policy π . It can be defined recursively by the formula

$$f^\pi(s) = \sum_{s'} P(s'|s, \pi(s)) \cdot [R(s, \pi(s), s') + \gamma f^\pi(s')], \quad (1)$$

where $P(s'|s, \pi(s))$ is the probability that the next state will be s' given that the current state is s , and we take action $\pi(s)$.

Given a policy π , a reward function R , and the transition probability function P , it is possible to compute the value function f^π by solving the system of linear equations containing one equation of the form of equation 1 for each possible state s . The system of equations




Figure 11. An Example of a Simple Backup.

Each arc is labeled with the probability of making the transition. The resulting states are labeled with their associated immediate reward and their value.

```

procedure POLICYITERATION( $P, R$ )
let  $\pi$  be an arbitrary initial policy
repeat until  $\pi$  is unchanged
    perform simple backups to compute  $f^{\pi}$  for each state (Equation 2).
    update  $\pi$  for each state (Equation 3).
end
end POLICYITERATION
    
```

Figure 12. The Policy-Iteration Algorithm.

can be solved using standard methods, such as Gaussian elimination or Gauss-Seidel iteration, or it can be solved iteratively by converting the equation into an assignment statement:

$$f^{\pi}(s) := \sum_{s'} P(s'|s, \pi(s)) \cdot [R(s, \pi(s), s') + \gamma f^{\pi}(s')], \quad (2)$$

This assignment statement is called a *simple backup* because it can be viewed as taking the current estimated value(s) $f^{\pi}(s')$ and “backing them up” to compute a revised estimate for $f^{\pi}(s)$. An example is shown in figure 11. In state s , we perform action $a = \pi(s)$. There are three possible resulting states:

$s_1, s_2,$ and $s_3,$

with probabilities 0.8, 0.1, and 0.1, respectively. The immediate rewards for reaching these states are 0, 0, and 2, respectively, and the estimated values $f(s')$ of the states are 10, 5, and 0, respectively. Assuming $\gamma = 1$, the new estimated value of $f^{\pi}(s)$ is

$$\begin{aligned} f^{\pi}(s) &:= (0 + 0.8 \cdot 10) + (0 + 0.1 \cdot 5) \\ &\quad + (2 + 0.1 \cdot 0) \\ &:= 10.5. \end{aligned}$$

To compute the value of a policy, the simple

backups can be performed in any convenient order—even at random. The only requirement is that γ must be small enough so that the sum of the expected discounted rewards converges, and the simple backups must be performed until $f^{\pi}(s)$ converges for every state s .

Given a value function $f^{\pi}(s)$, it is possible to compute a new policy π' that is guaranteed to be at least as good π . This new policy is computed by performing a one-step lookahead search and choosing the action whose backed-up value is the largest:

$$\pi'(s) := \operatorname{argmax}_a \sum_{s'} P(s'|s, a) [R(s, a, s') + \gamma f^{\pi}(s')]. \quad (3)$$

In words, we consider each possible action a , compute its expected backed-up value, and define $\pi'(s)$ to be the action that gives the maximum backed-up value.

By alternately computing the value f^{π} of policy π and then updating π using equation 2, we can converge to the optimal policy π^* and its optimal value function f^* . Note that if we plug f^* into equation 2, the new policy is unchanged—it is still the optimal policy. This defines the algorithm known as *policy iteration*, shown in figure 12. It is easy to show that policy iteration converges in a fixed number of iterations. Unfortunately, each iteration can be expensive because it requires computing the value of the current policy.

An alternative to policy iteration is to work directly with the value function. Bellman proved that the value function f^* of the optimal policy π^* is the unique fixed point of the Bellman equation:

$$f(s) := \max_a \sum_{s'} P(s'|s, a) [R(s, a, s') + \gamma f(s')] \quad (4)$$

In other words, we perform a one-step lookahead just as we did for the policy-improvement step of policy iteration, but instead of remembering the best action, we update our estimate of the value of state s . This is called a *Bellman backup*. It is more expensive than a simple backup because we must consider each possible action a and then each possible resulting state s' . By performing enough Bellman backups in every state, we can converge to the optimal value function f^* . This is called the value-iteration algorithm, and it is summarized in figure 13.

Unfortunately, value iteration can be more difficult than policy iteration because (1) each backup is a more expensive Bellman backup rather than a simple backup and (2) the value function can take a long time to converge. Indeed, it is possible for the optimal policy to have converged long before the value function converges.

A hybrid algorithm that combines aspects of

both value iteration and policy iteration is called *modified policy iteration*. This algorithm is essentially the same as policy iteration except that only a fixed number of simple backups are performed in each state in each iteration. Thus, the estimated value function for the current policy f^π does not completely converge to the correct value function, but under fairly mild conditions, it can be shown that the algorithm will still converge to the optimal policy.

All three of these algorithms—policy iteration, value iteration, and modified policy iteration—require performing backups in every state; so, their running time scales with the number of states. In fact, value iteration can run for an infinite amount of time without converging. Policy iteration requires time at least $O(n^3)$ for problems with n states just to compute the value of the policy in each iteration. For small problems, it is not a difficulty, but for problems of interest in AI, the state space often has 10^{20} or 10^{40} possible states, which renders these algorithms infeasible. Bellman termed this the *curse of dimensionality* because the number of states, and, hence, the running time, increases exponentially with the number of dimensions in the state space.

Another drawback of these algorithms is that they require a complete model of the system, by which I mean the transition probabilities $P(s'|s, a)$ and the reward function $R(s, a, s')$. There are many applications where this model is unavailable (for example, in a robot interacting with an unknown environment) or where the model cannot easily be converted into a transition probability matrix. For example, in the elevator control problem studied by Crites and Barto (1995), a software simulator is available that can take a current state s and a proposed action a and generate the next state s' according to the transition probability distribution. However, this distribution is not explicitly represented anywhere, so it is not available for direct use by a dynamic programming algorithm. The problem of constructing an explicit probability transition matrix and reward function has been called the *curse of modeling*, and in many problems, it is just as severe as the curse of dimensionality. Reinforcement learning algorithms provide a way of overcoming these two curses.

Reinforcement learning algorithms have introduced three key innovations: (1) stochastic approximation of backups, (2) value-function approximation, and (3) model-free learning. I discuss these innovations in the context of an algorithm known as TD(λ) developed by Sutton (1988).

```

procedure VALUEITERATION( $P, R$ )
  let  $f$  be an arbitrary initial value function
  repeat until  $f$  is unchanged in all states
    for each state  $s$ , perform a Bellman backup (Equation 4)
  end
  for each state  $s$ , compute the optimal policy (Equation 3)
end VALUEITERATION

```

Figure 13. The Value-Iteration Algorithm.

Temporal Difference Learning and TD(λ)

I begin by describing a simplified version of Sutton's TD(λ) algorithm, called TD(0). TD(0) is a method for computing approximate simple backups online. Suppose we are in state s , and we follow the current policy by taking action $a = \pi(s)$. If we are interacting with a real external environment (or with a simulator), the environment makes a probabilistic transition to a new state s' and produces the immediate reward $R(s, a, s')$. The TD(0) algorithm observes this new state and reward and updates the value function as follows:

$$f^\pi(s) := (1 - \alpha) f^\pi(s) + \alpha[R(s, a, s') + \gamma f^\pi(s')],$$

where α is a *learning rate parameter*. Typical values for α are between 0.01 and 0.5. (Technically, the α values must shrink to 0 over time for TD(0) to converge.)

The basic idea of TD(0) is that if we visit s many times and apply action a many times, then by sampling over time, we get the same effect as if we performed a simple backup. We can do this by sampling from the probability distribution $P(s'|s, a)$ rather than by having direct computational access to P .

This strategy allows TD(0) to compute the value of a policy without having an explicit model. The environment serves as its own model!

When Sutton developed TD(0), he also introduced a second idea. Instead of storing a separate value $f^\pi(s)$ for each state s , suppose we represent the value function as a neural network (or some other differentiable function approximator) of the form $f(s, W)$, where W is a vector of adjustable weights. With this representation, we can't directly assign a value to a state, but we can adjust the weights so that $f(s, W)$ is closer to the desired value by defining an error function:

$$J(W) = \frac{1}{2} (f^\pi(s, W) - [R(s, a, s') + \gamma f^\pi(s', W)])^2$$

This is (one half) the squared difference between the current estimated value of state s , $f^\pi(s, W)$, and the backed-up value, which is called the *temporal difference error*. Our goal is to




Figure 14. A Sequence of States.

The eligibility of each state (with $\lambda = 0.8$) is shown as a vertical bar.

modify W to reduce the temporal difference error $J(W)$. Differentiating $J(W)$ (and treating only the first occurrence of W as adjustable), we obtain the learning rule

$$W := W - \alpha \nabla_w f(s, W) (f^c(s, W) - [R(s, a, s') + \gamma f^c(s', W)]),$$

where $\nabla_w f(s, W)$ is the gradient of f with respect to the weights W . This takes a step of size α in the direction of the decreasing gradient scaled by the size of the temporal difference error.

By using a smoothly parameterized function approximation $f(s, W)$, TD(0) can circumvent the curse of dimensionality, provided that $f(s, W)$ can accurately approximate the true value function $f(s)$ with a small number of parameters W .

Sutton introduced one other wonderful idea in the TD(λ) algorithm—the *eligibility trace*. Suppose we have visited a sequence of states $s_1, s_2, \dots, s_t, s_{t+1}$, and we are updating the value of $f(s_t, W)$. Sutton suggested that we might want to update the values of the preceding states as well because the key idea of dynamic programming is to propagate information about expected rewards backward through the state space. Sutton proposed that we should remember the gradient $\nabla_{W_{t,i}} f(s_{t,i}, W_{t,i})$ for each state $s_{t,i}$ for $i = 1, \dots, n$, that we have visited. When we update $f(s_{t,i}, W_{t,i})$ by taking a step in the direction of $-\nabla_{W_{t,i}} f(s_{t,i}, W_{t,i})$, we also take a smaller step in the direction of $-\nabla_{W_{t,i-1}} f(s_{t,i-1}, W_{t,i-1})$ and an even smaller step in the direction of $-\nabla_{W_{t,i-2}} f(s_{t,i-2}, W_{t,i-2})$, and so on. Each step size will be decreased by a factor of $\lambda < 1$, which gives us the learning rule

$$W := W - \alpha \left(\sum_{i=0}^{\infty} \lambda^i \nabla_{W_{t-i}} f(s_{t-i}, W_{t-i}) \right) (f^\pi(s_t, W) - [R(s_t, a, s_{t+1}) + \gamma f^\pi(s_{t+1}, W)]).$$

The value of λ is called the *eligibility* of state $s_{t,i}$. Figure 14 shows the eligibility of a sequence of states as a bar graph.

The infinite sum

$$\sum_{i=0}^{\infty} \lambda^i \nabla_{W_{t-i}} f(s_{t-i}, W_{t-i})$$

can be implemented by maintaining a current gradient vector G :

$$G_t := \lambda G_{t-1} + \nabla_w f(s_t, W_t).$$

With this change, we get the full algorithm TD(λ) shown in figure 15. Readers familiar with the momentum method for stabilizing back propagation will note that the eligibility trace mechanism is similar. However, in the momentum method, previous weight changes are remembered, but in the eligibility trace mechanism, previous gradient vectors are remembered, and future temporal differences determine the step size along these previous gradients.

There have been many theoretical studies of the behavior of TD(λ). The most general results have been obtained by Tsitsiklis and Van Roy (1996). They analyze the case where the function approximator $f(s, W)$ is a linear combination of fixed (and arbitrary) orthogonal basis functions. To analyze how well $f(s, W)$ can approximate f^c , some notion of the distance between two value functions is needed. Tsitsiklis and Van Roy define the following measure:

$$\|f - f^c\|_D = (\sum_s [f(s) - f^c(s)]^2 D(s))^{1/2},$$

where $D(s)$ is the probability that the policy π visits state s .

With this measure, let

$$\hat{f}^\pi$$

be the best approximation of f^c that can be represented by a linear combination of the given basis functions. Tsitsiklis and Van Roy prove that TD(λ) will converge to a function f such that

$$\|f - f^\pi\|_D \leq \frac{\|\hat{f}^\pi - f^c\|_D}{1 - \gamma(1 - \lambda) / (1 - \lambda\gamma)}.$$

The quantity on the left-hand side is the error between the function f learned by TD(λ) and the true value function f^c for policy π . The numerator on the right-hand side is the inherent

approximation error resulting from the use of a linear combination of the given, fixed-basis functions. The denominator is the error that results from the fact that approximation errors in one state will be propagated backward to earlier states. For large λ , the denominator approaches 1 (no error), but for $\lambda = 0$, the denominator becomes $1 - \gamma$, which could be small and, hence, produce large errors. For this result to hold, it is essential that the backups performed by $\text{TD}(\lambda)$ be performed according to the current policy π . If this condition is not observed, then $\text{TD}(\lambda)$ may fail to converge.

The $\text{TD}(\lambda)$ algorithm provides a way of computing the value of a fixed policy without direct access to the transition probabilities and reward function. However, this algorithm is only of limited utility unless we can perform the policy-improvement step from equation 3, thereby implementing the policy-iteration algorithm. Unfortunately, policy improvement requires access to a model that can generate the possible next states and their probabilities.

Applications of $\text{TD}(\lambda)$

There are many domains where such a model is available. For example, in game-playing settings, such as backgammon, it is easy to compute the set of available moves from each state and the probabilities of all successor states. Hence, $\text{TD}(\lambda)$ can be combined with policy improvement to learn an approximately optimal policy for backgammon, which is what Tesauro (1995, 1992) did in his famous TD-GAMMON system.

TD-GAMMON uses a neural network representation of the value of a state. The state of the backgammon game is described by a vector of 198 features that encode the locations of the pieces on the board and the values shown on the dice. TD-GAMMON begins with a randomly initialized neural network. It plays a series of games against itself. At each step, it makes a full one-step lookahead search, applies the neural network to evaluate each of the resulting states, and makes the move corresponding to the highest backed-up value. In short, it applies equation 3 to compute the action to perform next, which is a form of local policy improvement. After making the move, it observes the resulting state and applies the $\text{TD}(\lambda)$ rule to update the value function. In effect, TD-GAMMON is executing a form of modified policy iteration where it alternates between one step of policy evaluation (the $\text{TD}(\lambda)$ update) and one step of policy improvement (the computation of the best move to make).

There is no guarantee that this algorithm will converge to the optimal policy. Indeed, it

```

procedure TD( $\pi, \gamma, \lambda, \gamma, \alpha, s_0$ )
  initialize  $G = 0$ 
  initialize  $W$  randomly
   $s_0$  is the starting state
  while  $W$  has not converged do
    take action  $a = \pi(s_t)$ 
    observe resulting state  $s_{t+1}$  and reward  $R(s_t, a, s_{t+1})$ 
     $G := \lambda G + \nabla_W f(s_t, W)$ 
     $W := W - \alpha G \cdot (f(s_t, W) - [R(s_t, a, s_{t+1}) + \gamma f(s_{t+1}, W)])$ 
  end while
  return  $W$  which defines  $f^{\pi, \gamma}$ 
end TD

```

Figure 15. The $\text{TD}(\lambda)$ Algorithm for Computing the Value of a Policy π .

is easy to construct examples where the strategy of always performing the best action based on the current approximation to the value function, $f(s, W)$, leads to a local minimum. For example, consider the navigation problem shown in figure 16. There is a large mountain separating the start state from the goal. A network of roads passes to the north of the mountain, and a similar (and shorter) network passes to the south. Suppose that our initial policy takes us to the north side, and we eventually reach the goal. After updating our value function using $\text{TD}(\lambda)$, suppose that the north side still appears to be shorter than the south side (because our estimates for the values of the states along the south side are too high). Then, in future trials, we will continue to take the north roads, and we will never try the southern route.

This is called the *problem of exploration*. The heart of the problem is that to find the optimal policy, it is necessary to prove that every off-policy action leads to expected results that are worse than the actions of the optimal policy. In this situation, it is essential to explore the southern path to determine whether it is worse (or better!) than the northern path. The strategy of always taking the action that appears to be optimal based on the current value function is called the *pure exploitation strategy*. The example in figure 16 shows that the pure exploitation strategy does not always find the optimal policy. Hence, online reinforcement learning algorithms must balance exploitation with exploration.

Fortunately, in backgammon, the random dice rolls inject so much randomness into the game that TD-GAMMON thoroughly investigates the possible moves in the game. Experimentally, the performance of TD-GAMMON is outstanding. It plays much better than any other computer program, and it is nearly as good as the world's best players. The results of three versions of the program in three separate matches




Figure 17. A Comparison of the Value Function f and the Q Function.

Black nodes represent situations where the agent has chosen an action. White nodes are states where the agent has not yet chosen an action.

The second approach is a model-free algorithm called Q -LEARNING, developed by Watkins (Watkins 1989; Watkins and Dayan 1992), which is the subject of the next section.

Model-Free Reinforcement Learning (Q -LEARNING)

Q -LEARNING is an online approximation of value iteration. The key to Q -LEARNING is to replace the value function $f(s)$ with an action-value function, $Q(s, a)$. The quantity $Q(s, a)$ gives the expected cumulative discounted reward of performing action a in state s and then pursuing the current policy thereafter. Hence, the value of a state is the maximum of the Q values for that state:

$$f(s) = \max_a Q(s, a).$$

We can write down the Q version of the Bellman equation as follows:

$$Q(s, a) = \sum_{s'} P(s'/s, a) [R(s, a, s') + \max_a \gamma Q(s', a)].$$

The role of the Q function is illustrated in figure 17, where it is contrasted with the value function f . In the left part of the figure, we see that the Bellman backup updates the value of $f(s)$ by considering the values $f(s')$ of states that result from different possible actions. In the right part of the figure, the analogous backup works by taking the best of the values $Q(s', a')$ and backing them up to compute an updated value for $Q(s, a)$.

The Q function can be learned by an algorithm that exploits the same insight as TD(0)—online sampling of the transition probabilities—and an additional idea—online sampling of the available actions. Specifically, suppose we have visited state s , performed action

a , and observed the resulting state s' and immediate reward $R(s, a, s')$. We can update the Q function as follows:

$$Q(s, a) := (1 - \alpha)Q(s, a) + \alpha[R(s, a, s') + \gamma \max_a Q(s', a)].$$

Suppose that every time we visit state s , we choose the action a uniformly at random. Then, the effect is to approximate a full Bellman backup (see equation 4). Each value $Q(s, a)$ is the expected cumulative discounted reward of executing action a in s , and the maximum of these values is $f(s)$. The random choice of a ensures that we learn Q values for every action available in state s . The online updates ensure that we experience the resulting states s' in proportion to their probabilities $P(s'/s, a)$.

In general, we can choose which actions to perform in any way we like, as long as there is a nonzero probability of performing every action in every state. Hence, a reasonable strategy is to choose the best action (that is, the one whose estimated value $Q(s, a)$ is largest) most of the time but to choose a random action with some small probability. Watkins proves that as long as every action is performed in every state infinitely many times, the Q function converges to the optimal function Q^* with probability 1.

Once we have learned Q^* , we must convert it into a policy. Although this conversion presented an insurmountable difficulty for TD(0), it is trivial for the Q representation. The optimal policy can be computed as

$$\pi^*(s) = \operatorname{argmax}_a Q^*(s, a).$$

Crites and Barto (1995) applied Q -LEARNING to a problem of controlling 4 elevators in a 10-story office building during peak “down traf-




Figure 18. Comparison of Learned Elevator Policy Q with Eight Published Heuristic Policies.

fic.” People press elevator call buttons on the various floors of the building to call the elevator to the floor. Once inside the elevator, they can also press destination buttons to request that the elevator stop at various floors. The elevator control decisions are (1) to decide, after stopping at a floor, which direction to go next and (2) to decide, when approaching a floor, whether to stop at the floor or skip it. Crites and Barto applied rules to make the first decision; so, the reinforcement learning problem is to learn whether to stop or skip floors. The goal of the controller is to minimize the square of the time that passengers must wait for the elevator to arrive after pressing the call button.

Crites and Barto used a team of four Q -learners, one for each of the four elevator cars. Each Q -function was represented as a neural network with 47 input features, 20 sigmoidal hidden units, and 2 linear output units (to represent $Q(s, stop)$ and $Q(s, skip)$). The immediate reward was the (negative of the) squared wait time since the previous action. They employed a form of random exploration in which exploratory actions are more likely to be chosen if they have higher estimated Q values.

Figure 18 compares the performance of the learned policy to that of eight heuristic algorithms, including the best nonproprietary algorithms. The left-hand graph shows the squared wait time, and the right-hand graph shows the percentage of passengers that had to wait more than 60 seconds for the elevator. The learned Q policy performs better than all the other methods.

Another interesting application of Q -LEARN-

ING is to the problem of assigning radio channels for cellular telephone traffic. Singh and Bertsekas (1997) showed that Q -LEARNING could find a much better policy than some sophisticated and complex published methods.

Open Problems in Reinforcement Learning

Many important problems remain unsolved in reinforcement learning, which reflects the relative youth of the field. I discuss a few of these problems here.

First, the use of multilayer sigmoidal neural networks for value-function approximation has worked, but there is no reason to believe that such networks are well suited to reinforcement learning. First, they tend to forget episodes (both good and bad) unless they are retrained on the episodes frequently. Second, the need to make small gradient-descent steps makes learning slow, particularly in the early stages. An important open problem is to clarify what properties an ideal value-function approximator would possess and develop function approximators with these properties. Initial research suggests that value-function approximators should be local averagers that compute the value of a new state by interpolating among the values of previously visited states (Gordon 1995).

A second key problem is to develop reinforcement methods for hierarchical problem solving. For very large search spaces, where the distance to the goal and the branching factor are big, no search method can work well. Often such large search spaces have a hierarchical

(or approximately hierarchical) structure that can be exploited to reduce the cost of search. There have been several studies of ideas for hierarchical reinforcement learning (for example, Dayan and Hinton [1993], Kaelbling [1993], and Singh [1992]).

The third key problem is to develop intelligent exploration methods. Weak exploration methods that rely on random or biased random choice of actions cannot be expected to scale well to large, complex spaces. A property of the successful applications shown previously (particularly, backgammon and job-shop scheduling) is that even random search reaches a goal state and receives a reward. In domains where success is contingent on a long sequence of successful choices, random search has a low probability of receiving any reward. More intelligent search methods, such as means-ends analysis, need to be integrated into reinforcement learning systems as they have been integrated into other learning architectures such as SOAR (Laird, Newell, and Rosenbloom 1987) and PRODIGY (Minton et al. 1989).

A fourth problem is that optimizing cumulative discounted reward is not always appropriate. In problems where the system needs to operate continuously, a better goal is to maximize the average reward per unit time. However, algorithms for this criterion are more complex and not as well behaved. Several new methods have been put forward recently (Mahadevan 1996; Ok and Tadepalli 1996; Schwartz 1993).

The fifth, and perhaps most difficult, problem is that existing reinforcement learning algorithms assume that the entire state of the environment is visible at each time step. This assumption is not true in many applications, such as robot navigation or factory control, where the available sensors provide only partial information about the environment. A few algorithms for the solution of hidden-state reinforcement learning problems have been developed (Littman, Cassandra, and Kaelbling 1995; McCallum 1995; Parr and Russell 1995; Cassandra, Kaelbling, and Littman 1994). Exact solution appears to be difficult. The challenge is to find approximate methods that scale well to large hidden-state applications.

Despite these substantial open problems, reinforcement learning methods are already being applied to a wide range of industrial problems where traditional dynamic programming methods are infeasible. Researchers in the area are optimistic that reinforcement learning algorithms can solve many problems that have resisted solution by machine-learning methods in the past. Indeed, the general problem of

choosing actions to optimize expected utility is exactly the problem faced by general intelligent agents. Reinforcement learning provides one approach to attacking these problems.

Learning Stochastic Models

The final topic that I discuss is the area of learning stochastic models. Traditionally, researchers in machine learning have sought general-purpose learning algorithms—such as the decision tree, rule, neural network, and nearest-neighbor algorithms—that could efficiently search a large and flexible space of classifiers for a good fit to training data. Although these algorithms are general, they have a major drawback. In a practical problem where there is extensive prior knowledge, it can be difficult to incorporate this prior knowledge into these general algorithms. A secondary problem is that the classifiers constructed by these general learning algorithms are often difficult to interpret—their internal structure might not have any correspondence to the real-world process that is generating the training data.

Over the past five years or so, there has been tremendous interest in a more knowledge-based approach based on stochastic modeling. A *stochastic model* describes the real-world process by which the observed data are generated. Sometimes, the terms *generative stochastic model* and *causal model* are used to emphasize this perspective. The stochastic model is typically represented as a *probabilistic network*—a graph structure that captures the probabilistic dependencies (and independencies) among a set of random variables. Each node in the graph has an associated probability distribution, and from these individual distributions, the joint distribution of the observed data can be computed. To solve a learning problem, the programmer designs the structure of the graph and chooses the forms of the probability distributions, yielding a stochastic model with many free parameters (that is, the parameters of the node-probability distributions). Given a training sample, learning algorithms can be applied to determine the values of the free parameters, thereby fitting the model to the data. Once a stochastic model has been learned, probabilistic inference can be carried out to support tasks such as classification, diagnosis, and prediction.

More details on probabilistic networks are given in two recent textbooks: Jensen (1996) and Castillo, Gutierrez, and Hadi (1997).

Probabilistic Networks

Figure 19 is an example of a probabilistic net-




Figure 19. A Probabilistic Network for Diabetes Diagnosis.

work that might be used to diagnose diabetes. There are six variables (with their abbreviations): (1) Age: age of patient (*A*); (2) Preg: number of pregnancies (*N*); (3) Mass: body mass (*M*); (4) Insulin: Blood-insulin level (after a glucose-tolerance test) (*I*); (5) Glucose: blood-glucose level (after a glucose-tolerance test) (*G*); and (6) Diabetes: true if the patient has diabetes (*D*).

In a medical diagnosis setting, the first five variables would be observed and then the computer would estimate the probability that the patient has diabetes (that is, estimate the probability that the diabetes variable is true).

This network corresponds to the following decomposition of the joint-probability distribution among the six variables:

$$P(A, N, M, I, G, D) = P(A) \cdot P(N) \cdot P(M|A, N) \cdot P(D|M, A, N) \cdot P(I|D) \cdot P(G|I, D).$$

Each node in the network corresponds to a probability distribution of the form $P(\text{Node}|\text{Parents})$, where the *Parents* of *Node* are the nodes with arcs pointing to *Node*. In other words, if we believe the network is a correct representation of the relationships among the variables, then it should be possible to factor the joint probability distribution into the product of these smaller distributions.

The structure of the network—particularly the arcs that are absent—can be viewed as specifying conditional independencies. Two variables *A* and *B* are conditionally independent given *C* if

$$P(A, B|C) = P(A|C) \cdot P(B|C).$$

In the graph, Age affects Insulin only through the Diabetes node; so, Age and Insulin are conditionally independent given Diabetes. More generally, given the values of its parents, a node is independent of all other nodes in the graph except its descendants. Formally,

$$P(A, B|\text{Parents}(A)) = P(B|\text{Parents}(A)) \cdot P(A|\text{Parents}(A)),$$

unless *B* is a descendant of *A*.

In addition to specifying the structure of the network, we need to specify how each probability distribution is represented. One standard approach is to discretize the variables into a small number of values and represent each probability distribution as a table. For example, suppose we discretized Age into the values {0–25, 26–50, 51–75, > 75}, Preg into the values {0, 1, > 1}, and Mass into the values {0–50 kilograms (kg), 51–100 kg, > 100 kg}. Then, the probability distributions for these three nodes could be represented by the probability tables shown in table 3. The learning task is to fill in the probability values in these tables. The table for $P(A)$ requires three independent parameters (because the four values must sum to 1). The table for $P(N)$ requires 2 parameters, and the table for $P(M|A, N)$ requires 24 parameters (because each row must sum to 1). Similar tables would be required for the other three nodes in the network.

Given a set of training examples, this learning problem is easy to solve. Each probability can be computed directly from the training data. For example, the cell $P(N=1)$ can be computed as the proportion of patients in the sample that had exactly 1 pregnancy. The parameter $P(M=51-100|A=26-50, N=1)$ is the fraction of training examples with Age = 26–50 and Preg = 1 that have a Mass = 51–100 kg. Technically, these are the maximum-likelihood estimates of each of the probabilities. A difficulty that can arise is that some of the cells in the tables can have few examples; so, the resulting probability estimates are uncertain. One solution is to smooth probabilities for adjacent cells in the table. For example, we might require that $P(M=51-100|A=26-50, N=1)$ have a value similar to $P(M=51-100|A=26-50, N=>1)$. Of course, we could also take an ensemble approach and generate an ensemble of fitted stochastic models, as described in the section Ensembles of Classifiers.

The process of learning a stochastic model consists of three steps: (1) choosing the graph structure, (2) specifying the form of the probability distribution at each node in the graph, and (3) fitting the parameters of these probability distributions to the training data. In

Age	$P(A)$		Preg	$P(N)$
0-25			0	
26-50			1	
51-75			>1	
> 75				

		$P(M/A, N)$		
Age	Preg	0-50	51-100	>100
0-25	0			
0-25	1			
0-25	> 1			
26-50	0			
26-50	1			
26-50	>1			
51-75	0			
51-75	1			
51-75	>1			
>75	0			
>75	1			
>75	> 1			

Table 3. Probability Tables for the Age, Preg, and Mass Nodes from Figure 19.

A learning algorithm must fill in the actual probability values based on the observed training data.

most current applications, steps 1 and 2 are performed by a user, and step 3 is performed by a learning algorithm. However, later in this section, I briefly discuss methods for automating step 1—learning the graph structure.

Once we learn the model, how can we apply it to predict whether new patients have diabetes? For a new case, we observe the values of all the variables in the model except for the Diabetes node. Our goal is to compute the probability of the node; that is, we seek the distribution $P(D/A, N, M, I, G)$. We could precompute this distribution offline before observing any of the data, but the resulting conditional probability table would be immense. A better approach is to wait until the values of the variables have been observed and then compute the single corresponding row of the class probability table.

This inference problem has been studied intensively, and a general and elegant algorithm—the junction-tree algorithm—has been developed (Jensen, Lauritzen, and Olesen 1990). In addition, efficient online algorithms have been discovered (D'Ambrosio 1993). In the

worst case, these algorithms require exponential time, but if the probabilistic network is sparsely connected, the running time is reasonable.

The Naive Bayes Classifier

A simple approach to stochastic modeling for classification problems is the so-called naive Bayes classifier (Duda and Hart 1973). In this approach, the training examples are assumed to be produced by the probabilistic network shown in figure 20, where the class variable is y , and the features are x_1, \dots, x_n . According to this model, the environment generates an example by first choosing (stochastically) which class to generate. Then, once the class is chosen, the features describing the example are generated independently according to their individual distributions $P(x_i|y)$.

In most applications, the values of the features are discretized so that each feature takes on only a small number of discrete values. The probability distributions are represented as tables as in the diabetes example, and the network can be learned directly from the training data by counting the fraction of examples in




Figure 20. Probabilistic Network for the Naive Bayes Classifier.

each class that take on each feature value.

Because of the simple form of the network, it is easy to derive a classification rule through the application of Bayes' rule. Suppose there are only two classes, 1 and 2. Then, our decision rule is to classify a new example into class 1 if $P(y = 1|\mathbf{x}) > P(y = 2|\mathbf{x})$ or, equivalently, if $P(y = 1|\mathbf{x})/P(y = 2|\mathbf{x}) > 1$. By Bayes's rule, we can write

$$\begin{aligned} P(y = 1|\mathbf{x}) &= P(\mathbf{x}|y = 1) \cdot P(y = 1)/P(\mathbf{x}) \\ P(y = 2|\mathbf{x}) &= P(\mathbf{x}|y = 2) \cdot P(y = 2)/P(\mathbf{x}). \end{aligned}$$

Dividing the first equation by the second allows us to cancel the normalizing denominator $P(\mathbf{x})$ and obtain

$$\frac{P(y = 1|\mathbf{x})}{P(y = 2|\mathbf{x})} = \frac{P(\mathbf{x}|y = 1) \cdot P(y = 1)}{P(\mathbf{x}|y = 2) \cdot P(y = 2)}.$$

The quantity $P(\mathbf{x}|y = 1)$ is just the product of the individual probabilities $P(x_i|y = 1)$; so, we have

$$\begin{aligned} \frac{P(y = 1|\mathbf{x})}{P(y = 2|\mathbf{x})} &= \frac{P(x_1|y = 1) \cdot P(x_2|y = 1)}{P(x_1|y = 2) \cdot P(x_2|y = 2)} \\ &\quad \cdots \frac{P(x_n|y = 1) \cdot P(y = 1)}{\cdots P(x_n|y = 2) \cdot P(y = 2)}. \end{aligned}$$

This gives the decision rule that we should classify an example into class 1 if and only if

$$\begin{aligned} \frac{P(x_1|y = 1)}{P(x_1|y = 2)} \cdot \frac{P(x_2|y = 1)}{P(x_2|y = 2)} \cdots \\ \frac{P(x_n|y = 1)}{P(x_n|y = 2)} \cdot \frac{P(y = 1)}{P(y = 2)} > 1. \end{aligned}$$

Despite the fact that the naive Bayes model barely deserves the name *model* in many applications, it performs surprisingly well. Figure 21 compares the performance of *c4.5* to the naive Bayes classifier on 28 benchmark tasks (Domingos and Pazzani 1996). The results show that except for a few domains where naive Bayes performs badly, it is typically com-

petitive with, or superior to, *c4.5*. Domingos and Pazzani showed that the algorithm is robust to violations of the assumption that the features are generated independently.

Naive Unsupervised Learning

An important application of stochastic models is to problems of unsupervised learning. In unsupervised learning, we are given a collection of examples $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, and our goal is to construct some model of how these examples are generated. For example, we might believe that these examples belong to some collection of classes, and we want to determine the properties of these classes. This is sometimes called *clustering* the data into classes, and many algorithms have been developed that apply a measure of the distance between two examples to group together nearby examples.

A stochastic modeling approach for unsupervised learning works essentially the same way as for supervised learning. We begin by defining the structure of the model as a probabilistic network. One commonly used model is the same naive network used by the naive Bayes classifier shown in figure 20.

Although we can use the same network structure, the problem of learning the parameters of the network is much more difficult because our data do not contain the values of the class variable y . This is a simple case of the problem of fitting stochastic models that contain *hidden variables*, which are variables whose values are not observed in the training data.

Many different algorithms have been developed for fitting networks containing hidden variables. As with the naive Bayes classification algorithm, the basic goal (at least for this article) is to compute the maximum-likelihood estimates of the parameters of the network, which I shall refer to as the vector of weights W . In other words, we want to find the value of W that maximizes $P(S|W)$, where S is the observed training sample. This is typically formulated as the equivalent problem of maximizing the log likelihood: $\log P(S|W)$. Under the assumption that each training example in S is generated independently, this is equivalent to maximizing $\sum_i \log P(\mathbf{x}_i|W)$, where \mathbf{x}_i is the i -th training example.

I briefly sketch three algorithms: (1) gradient descent, (2) expectation maximization, and (3) Gibbs sampling.

Gradient Descent for Bayes Networks

Russell et al. (1995) describe a method for computing the gradient of the log likelihood: $\nabla_w \log P(\mathbf{x}_i|W)$. Let us focus on a particular node in the network representing variable V . Let U be the parents of V , and let $w = P(V = v|U = u)$

be the entry in the conditional probability table for V when $V = v$ and $U = u$. Then, Russell et al. show that

$$\frac{\partial \log P(\mathbf{x}_i | W)}{\partial w} = \frac{P(V = v, U = u | W, \mathbf{x}_i)}{w}.$$

The conditional probability in the numerator can be computed by any algorithm for inference in probabilistic networks (including the junction-tree algorithm mentioned earlier).

With this formula, the gradient of the parameters of a network can be computed with respect to a training sample S . We can then apply standard gradient-descent algorithms, such as fixed step-size methods or the conjugate gradient algorithm, to search for the maximum-likelihood set of weights. One subtlety is that we must ensure that the weights lie between 0 and 1 and sum to 1 appropriately. It suffices to constrain and renormalize them after each step of gradient descent. Russell et al. have tested this algorithm on a wide variety of probabilistic network structures.

The Expectation Maximization Algorithm The second algorithm is the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1976). EM can be applied to probabilistic networks if the node-probability distributions belong to the exponential family of distributions (which includes the binomial, multinomial, exponential, poisson, and normal distributions). In particular, conditional probability tables have multinomial distributions; so, the EM algorithm can be applied to the case described in this article.

The EM algorithm is an iterative algorithm that starts with an initial value for W and incrementally modifies W to increase the likelihood of the observed data. One way to understand EM is to imagine that each training example is augmented to include parameters describing values of the hidden variables. For concreteness, consider the simple case of figure 20 with the class variable y hidden. Assume y takes on two values, 1 and 2. Then, we would augment each training example with $P(y = 1 | \mathbf{x})$ and $P(y = 2 | \mathbf{x})$. (Actually, the $P(y = 2 | \mathbf{x})$ is redundant in this case because $P(y = 2 | \mathbf{x}) = 1 - P(y = 1 | \mathbf{x})$). More generally, each observed example \mathbf{x}_i will be augmented with the expected values of the sufficient statistics for describing the probabilistic distribution of the hidden variables.

The EM algorithm alternates between two steps until W converges:

First is the *E-step*: Given the current value of W , compute the augmentations $P(y_i = 1 | \mathbf{x}_i)$ and $P(y_i = 2 | \mathbf{x}_i)$ for each example \mathbf{x}_i .

Second is the *M-step*: Given the augmented data set, compute the maximum-likelihood estimates for W under the assumption that the




Figure 21. Comparison of *c4.5* and the Naive Bayesian Classifier on 28 Data Sets.

probability distribution of the values of the hidden variables is correctly specified in each augmented training example.

In the case of figure 20, the E-step applies Bayes's theorem:

$$P(y_i = 1 | \mathbf{x}_i) = P(\mathbf{x}_i | y_i = 1) \cdot P(y_i = 1) / P(\mathbf{x}_i).$$

All the quantities on the right-hand side can be computed given the current value for W . The quantity $P(\mathbf{x}_i | y_i = 1)$ is the product of $P(x_j | y_i = 1)$ for each feature j , and $P(y_i = 1)$ is the current estimated probability of generating an example in class 1, $P(y = 1)$.

The M-step for naive Bayes classification must estimate each of the weights from the augmented training examples. To estimate $P(y = 1)$, we sum the augmented value $P(y_i = 1)$ over all i and divide by the sample size m :

$$P(y = 1) = \frac{1}{m} \sum_i P(y_i = 1).$$

To estimate the conditional probability that feature x_j is 1 for examples in class 1, we take each training example i that has $x_j = 1$, sum up the augmented values $P(y_i = 1)$, and divide by the total $P(y = 1)$:

$$P(x_j = 1 | y = 1) = \frac{1}{P(y = 1)} \sum_{\{i | x_j = 1\}} P(y_i = 1)$$

In effect, we treat each augmented training example as if it were a member of class 1 with probability $P(y_i = 1)$ and as if it were a member of class 2 with probability $1 - P(y_i = 1)$.

A well-known application of the EM algorithm in unsupervised clustering is the AUTOCLASS program (Cheeseman et al. 1988). In addition to discrete variables (of the kind I have been discussing), AUTOCLASS can handle contin-

uous variables. Instead of computing the maximum-likelihood estimates of the parameters, it adopts a prior probability distribution over the parameter values and computes the maximum a posteriori probability (MAP) values. This is easily accomplished by a minor modification of EM. One of the most interesting applications of AUTOCLASS was to the problem of analyzing the infrared spectra of stars. AUTOCLASS discovered a new class of star, and this discovery was subsequently accepted by astronomers (Cheeseman et al. 1988).

Gibbs Sampling The final algorithm that I discuss is a Monte Carlo technique called *Gibbs sampling* (Geman and Geman 1984). Gibbs sampling is a method for generating random samples from a joint-probability distribution $P(A_1, \dots, A_n)$ when sampling directly from the joint distribution is difficult. Suppose we know the conditional distribution of each variable A_i in terms of all the others: $P(A_1|A_2, \dots, A_n)$, $P(A_2|A_1, A_3, \dots, A_n)$, ..., $P(A_n|A_1, \dots, A_{n-1})$. The Gibbs sampler works as follows: We start with a set of arbitrary values, a_1, \dots, a_n , for the random variables. For each value of i , we then sample a new value a_i for random variable A_i according to the distribution $P(A_i|A_1 = a_1, \dots, A_{i-1} = a_{i-1}, A_{i+1} = a_{i+1}, \dots, A_n = a_n)$. If we repeat this long enough, then under certain mild conditions the empirical distribution of these generated points will converge to the joint distribution.

How is this useful for learning in probabilistic networks? Suppose we want to take a full Bayesian approach to learning the unknown parameters in the network. In such cases, we want to compute the posterior probability of the unknown parameters given the data. Let W denote the vector of unknown parameters. In a full Bayesian approach, we treat W as a random variable with a prior probability distribution $P(W)$. Given the training examples $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$, we want to compute $P(W|\mathbf{x}_1, \dots, \mathbf{x}_m)$, which is difficult because of the hidden classes y_1, \dots, y_m . However, using the Gibbs sampler, we can generate samples from the distribution

$$P(W, y_1, \dots, y_m | \mathbf{x}_1, \dots, \mathbf{x}_m).$$

Then, by simply ignoring the y values, we obtain samples for W : W_1, \dots, W_ℓ .

These W values constitute an ensemble of learned values for the network parameters. To classify a new data point \mathbf{x} , we can apply each of the values of W_ℓ to predict the class of \mathbf{x} and have them vote. The voting is with equal weight. If some W values have higher posterior probability than others, then they will appear more often in our sample.

To apply the Gibbs sampler to our unsupervised learning problem, we begin by choosing random initial values for the parameters W of

the network and setting ℓ to zero. We then repeat the following loop L times.

First, compute new values for y_1, \dots, y_m : From the probabilistic network in figure 20, we can compute $P(y_j|W, \mathbf{x})$ because we know (current guesses for) W and (observed values for) \mathbf{x} . To sample from this distribution, we flip a biased coin with probability of heads $P(y_j|W, \mathbf{x})$.

Second, compute new values for W : Let w_0 be the parameter that represents the probability of generating an example from class 0. Suppose the prior distribution, $P(w_0)$, is the uniform distribution for all values $0 \leq w_0 \leq 1$. Let m_0 be the number of training examples (currently) assigned to class 0. Then, the posterior probability $P(w_0|y_1, \dots, y_m)$ has a special form known as a beta distribution with parameters $m_0 + 1$ and $m - m_0 + 1$. Algorithms are available for drawing samples from this distribution.

Similarly, let w_{j0} be the parameter that represents the probability that the j -th feature will have the value v when drawn from class 0: $P(x_j = v|y_j = 0)$. Again assuming uniform priors for $P(w_{j0})$, this variable also has a beta distribution with parameters $c_{j0} + 1$ and $m_0 - c_{j0} + 1$, where c_{j0} is the number of training examples in class 0 having $x_j = v$. As with w_0 , we can sample from this distribution.

We can do the same thing for the parameters concerning class 1: w_{j1} .

Third, record the value of the parameter vector. Let $\ell := \ell + 1$ and set $W_\ell := W$. To allow the Gibbs sampler to converge to a stationary distribution, we should perform some number of iterations before recording L values of the parameter vector. This procedure gives us an ensemble of L probabilistic networks that can be applied to classify new data points.

There is a close resemblance between Gibbs sampling and the EM algorithm. In both algorithms, the training examples are augmented with information about the hidden variables. In EM, this information describes the probability distribution of the hidden variables, whereas in Gibbs sampling, this information consists of random values drawn according to the probability distribution. In EM, the parameters are recomputed to be their maximum-likelihood estimates based on the current values of the hidden variables, whereas in Gibbs sampling, new parameter values are sampled from the posterior distribution given the current values of the hidden variables. Hence, EM can be viewed as a maximum-likelihood approximation (or MAP approximation) to Gibbs sampling.

One potential problem that can arise in Gibbs sampling is caused by symmetries in the stochastic model. In the case we are considering, for example, suppose there are two true

underlying hidden classes, class A and class B. We want the model to generate examples from class A when $y = 1$ and from class B when $y = 2$. However, there is nothing to force the model to do this. It could just as easily use $y = 1$ to represent examples of class B and $y = 2$ to represent examples of class A. If permitted to run long enough, in fact, the Gibbs sampler should explore both possibilities. If our sample of weight vectors W_i includes weight vectors corresponding to both of these alternatives, then when we combine these weight vectors, we will get bad results. In practice, this problem often does not arise, but in general, steps must be taken to remove symmetries from the model (see Neal [1993]).

Gibbs sampling is a general method; it can often be applied in situations where the EM algorithm cannot. The generality of Gibbs sampling has made it possible to construct a general-purpose programming environment, called BUGS, for learning stochastic models (Gilks, Thomas, and Spiegelhalter 1993). In this environment, the user specifies the graph structure of the stochastic model, the form of the probability distribution at each node, and prior distributions for each parameter. The system then develops a Gibbs sampling algorithm for fitting this model to the training data. BUGS can be downloaded from www.mrc-bsu.cam.ac.uk/bugs/.

More Sophisticated Stochastic Models

This section presents three examples of more sophisticated stochastic models that have been developed: (1) the hierarchical mixture of experts (HME) model, (2) the hidden Markov model (HMM), and (3) the dynamic probabilistic network (DPN). Like the naive model from figure 20, these models can be applied to a wide number of problems without performing the kind of detailed modeling of causal connections that we performed in the diabetes example from figure 19.

The Hierarchical Mixture of Experts

The HME model (Jordan and Jacobs 1994) is intended for supervised learning in situations where one believes the training data are being generated by a mixture of separate experts. For example, in a speech-recognition system, we might face the task of distinguishing the spoken words *bee*, *tree*, *gate*, and *mate*. This task naturally decomposes into two hard subproblems: (1) distinguishing *bee* from *tree* and (2) distinguishing *gate* from *mate*.

Figure 22 shows the probabilistic network for a simple mixture-of-experts model for this case. According to this model, a training example (\mathbf{x}_i, y_i) is generated by first generating the

data points \mathbf{x}_i , then choosing an expert e_i stochastically (depending on the value of \mathbf{x}_i), and then choosing the class y_i depending on the values of \mathbf{x}_i and e_i .

There are two things to note about this model. First, the direction of causality is reversed from the naive Bayes network of figure 20. Second, if we assume that all the features of each \mathbf{x}_i will always be observed, we do not need to model the probability distribution $P(X)$ on the bottom node in the graph. Third, unless we make some strong assumptions, the probability distribution $P(Y|X, E)$ is going to be extremely complex because it must specify the probability of the classes as a function of every possible combination of features for X and expert E .

In their development of this general model, Jordan and Jacobs assume that each probability distribution has a simple form (see figure 23, which shows the model as a kind of neural network classifier). Each value of the random variable E specifies a different expert. The input features \mathbf{x} are fed into each of the experts and into a gating network. The output of each expert is a probability distribution over the possible classes. The output of the gate is a probability distribution over the experts. This overall model has the following analytic form:

$$P(y | \mathbf{x}) = \sum_e g_e(\mathbf{x}) p_e(y | \mathbf{x}),$$

where the index e varies over the different experts. The value $g_e(\mathbf{x})$ is the output of the gating network for expert e . The value $p_e(y|\mathbf{x})$ is the probability distribution over the various classes output by expert e .

Jordan and Jacobs have investigated networks where the individual experts and the gating network have simple forms. In a two-class problem, each expert has the form

$$P_e(y | \mathbf{x}) = \sigma(w_e^T \mathbf{x}),$$

where

$$w_e^T$$

is the transpose of a vector of parameters, \mathbf{x} is the vector of input feature values, and σ is the usual logistic sigmoid function $1/(1 + \exp(\cdot))$.

The gating network was described earlier in Ensembles of Classifiers. The gating values are computed according to

$$z_e = v_e^T \mathbf{x}$$

$$g_e = e^{z_e} / \sum_u e^{z_u}.$$

In other words, z_e is the dot product of a weight vector v_e and the input features \mathbf{x} . The output g_e is the soft-max of the z_e values. The g_e values are all positive and sum to 1. This is known as the *multinomial logit model* in statistics.

The problem of learning the parameters for




Figure 22.
Probabilistic Model
Describing a Mixture
of Experts.

a mixture-of-experts model is similar to the problem of unsupervised learning, except that here, the hidden variable is not the class y of each training example i but, rather, the expert e_i that was responsible for generating the training example. If we knew which expert generated each training example, then we could fit the parameters directly from the training data, as we did for the naive Bayes algorithm. Jordan and Jacobs have applied both gradient-descent and the EM algorithm to solve this learning problem. For the particular choice of sigmoid and soft-max functions for the experts and gates, the EM algorithm has a particularly efficient implementation as a sequence of weighted least squares problems.

The simple one-level hierarchy of experts shown in figure 22 can be extended to deeper hierarchies. Figure 24 shows a three-level hierarchy. All the fitting algorithms apply to this more general case as well.

The Hidden Markov Model Figure 25 shows the probabilistic network of a hidden Markov model (HMM) (Rabiner 1989). An HMM generates examples that are strings of length n over some alphabet A of letters. Hence, each example \mathbf{x} is a string of letters: $o_1 o_2 \dots o_n$ (where the o stands for *observable*). To generate a string, the HMM begins by generating an initial state s_1 according to probability distribution $P(s_1)$. From this state, it then generates the first letter of the string according to the distribution $P(o_1/s_1)$. It then generates the next state s_2 according to $P(s_2/s_1)$. Then, it generates the second letter according to $P(o_2/s_2)$, and so on.

In some applications, the transition probability distribution is the same for all pairs of adjacent state variables: $P(s_i/s_{i-1})$. Similarly, the output (or emission) probability distribution

can be the same for all states: $P(o|s)$. In this variation, the HMM is equivalent to a stochastic finite-state automaton (FSA). At each time step, the FSA makes a probabilistic state transition and then generates an output letter.

Another common variation on HMMs is to include an absorbing state or halting state as one of the values of the state variable. If the HMM makes a transition into this state, it terminates the string being generated, permitting HMMs to model strings of variable length.

HMMs have been applied widely in speech recognition, where the alphabet of letters consists of frames of the speech signal (Rabiner 1989). Each word in the language can be modeled as an HMM. Given a new spoken word, a speech-recognition system computes the likelihood that each of the word HMMs generated this spoken word. The recognizer then predicts the most likely word. A similar analysis can be applied at the level of whole sentences by concatenating word-level HMMs. Then, the goal is to find the sentence most likely to have generated to speech signal.

To learn an HMM, a set of training examples is provided, where each example is a string. The sequence of states that generated the string is hidden. Once again, however, the EM algorithm can be applied. In the context of HMMs, it is known as the Baum-Welch, or forward-backward, algorithm. In the E-step, the algorithm augments each training example with statistics describing the hypothesized string of states that generated the example. In the M-step, the parameters of the probability distributions are reestimated from the augmented training data.

Stolcke and Omohundro (1994) developed algorithms for learning the structure and parameters of HMMs from training examples. They applied their techniques to several problems in speech recognition and incorporated their algorithm into a speaker-independent speech-recognition system.

The Dynamic Probabilistic Network In the HMM, the number of values for the state variable at each point in time can become large. Consequently, the number of parameters in the state-transition probability distribution $P(s_i/s_{i-1})$ can become intractably large. One solution is to represent the causal structure within each state by a stochastic model. For example, consider a mobile robot with a steerable television camera. The images observed by the camera will be the output alphabet of the HMM. The hidden state will consist of the location of the robot within a room and the direction the camera is pointing. Suppose there are 100 possible locations and 15 possible camera direc-




Figure 23. The Mixture-of-Experts Model Viewed
as a Specialized Neural Network.

tions. In an HMM, the hidden state will be a single variable with 1500 possible values.

However, suppose that the robot has separate commands to change its location and steer its camera. At each time step, it chooses to perform exactly one of these two actions. Then, it makes sense to represent the hidden state by two separate state variables: (1) robot location and (2) camera direction. Figure 26 shows the resulting stochastic model, which is variously called a DPN (Kanazawa, Koller, and Russell 1995), a dynamic belief network (Dean and Kanazawa 1989), and a factorial HMM (Ghamramani and Jordan 1996).

Unfortunately, inference and learning with DPNs is computationally challenging. The overall approach of applying the EM algorithm or Gibbs sampling is still sound. However, the E-step of computing the augmented training examples is itself difficult. Ghahramani and Jordan (1996) and Kanazawa, Koller, and Russell (1995) describe algorithms that can perform approximate E-steps. A recent review of this active research area can be found in Smyth, Heckerman, and Jordan (1997).

Application-Specific Stochastic Models

A major motivation for the stochastic modeling approach to machine learning is to communicate background knowledge to the learning algorithm. Although the general models discussed here achieve this goal to some extent, it is possible to go much further in this direction, and many applications of machine learning are pursuing this approach. To give a flavor of the kinds of model being developed, I describe one example from the many recently published papers.

Revw, Williams, and Hinton (1996) developed a stochastic model for handwritten digit recognition, shown in figure 27. To generate a digit according to this model, we first randomly choose one of the 10 digits, which determines the home locations of 8 points that control a uniform B spline (denoted h_1, \dots, h_8 in the figure). The B spline specifies the shape of the digit. The next step randomly perturbs these control points (using a Gaussian distribution) to produce the control points that will be used to generate the handwritten digit (denoted k_1, \dots, k_8 in the figure).

Next we generate a six-degree-of-freedom affine transformation that models translation, rotation, skewing, and so on. From the randomly chosen affine transformation and the control points, we deterministically lay out a sequence of beads along the spline curve. These beads act as circular Gaussian ink gener-




Figure 24. The Hierarchical Mixture-of-Experts Model.




Figure 25. The Probabilistic Network for a Hidden Markov Model.

ators that ink the pixels in the image. A parameter σ specifies the standard deviation of the Gaussian ink generators. The beads are placed a distance 2σ apart along the spline; hence, the number of beads varies as a function of σ . At this point, we have made all the choices that are independent of the image pixels.

Now, the model generates N inked pixels (indicated in the diagram by the box with an N in the lower left corner). To generate each pixel, it first generates a Boolean variable that indicates whether the pixel will be a noise pixel or a pixel from the digit. If it is a noise pixel, the location is chosen uniformly at random and inked. If it is a digit pixel, one of the I beads is chosen at random (specified by the variable bead index), and a pixel is inked according to a symmetric Gaussian probability distribution.

The parameters of the model include the probability of noise pixels, the standard devia-




Figure 26. A Simple Dynamic Probabilistic Network with Two Independent State Variables: Robot Location (L_i) and Camera Direction (S_i).




Figure 27. Probabilistic Network for Generating a Digit and Choosing N Pixels to Ink.

tion σ of the Gaussian ink generators, the variance of the Gaussian used to perturb the control points, the parameters controlling the affine transformation, and the home locations of the eight control points.

We want to perform two tasks with this model: (1) classification and (2) learning. Let us consider classification first. Given an image and 10 learned spline models, our goal is to find the spline model that is most likely to have generated the image. We do this by computing the probability that the spline model for each digit generated the image. This computation is difficult because we are given only the locations of the inked pixels and the (hypothesized) class of the digit. In principle, we should consider all possible locations for the control points, all possible values for σ , and all possible choices for whether each inked pixel is a noise pixel or a digit pixel. Each such combination could have generated the observed image (with a certain probability). We should integrate over these parameters to compute

the probability that a given digit model generated the given image.

In practice, Revow et al. take a maximum-likelihood approach. They compute the values for the control points (k_1, \dots, k_8), σ , and the noise-digit choices that maximize the likelihood of generating the observed image. The EM algorithm is applied to compute this. Figure 28 shows the fitting of the model for a 3 and for a 5 to a typical image. Notice that during the fitting, σ is initialized to a large value so that the ink generators are likely to capture inked pixels. When EM begins to converge, σ is decreased (according to a given schedule), and new beads are positioned along the spline using the current value of σ . Then, fitting is resumed. This task is performed approximately six times for each image.

Now let's consider learning. The goal of learning is to learn for each digit the home locations of the control points and the variance for perturbing these control points. In addition, the parameters controlling the affine transformation must be learned, but these parameters are assumed to be the same for all classes. To solve this learning problem, Revow et al. again apply EM. In the E-step, the training examples are augmented with the maximum-likelihood locations of the control points and the value of σ , which are computed using a nested EM, as described earlier. In the M-step, the average home location of each control point and the average σ are computed. The model for each digit is trained by fitting only to training examples of that digit. The running time is dominated by the cost of the inner (classification) EM algorithm.

Revow et al. report results that are competitive with the best-known methods on the task of recognizing digits from U.S. mail zip codes. To achieve this performance, they use some postprocessing steps that analyze how well each digit model fits the image. They also considered learning mixtures of digit models for cases where there are significantly different ways of writing a digit (for example, the digit 7 with and without a central horizontal bar).

The advantage of the stochastic modeling approach is that learning is fast both computationally, because EM is quick, and statistically, because there are only 22 parameters in each digit model (2 coordinates for each control point and 6 affine transformation parameters). Another advantage is that the digit images do not need to be preprocessed (for example, to remove slants and scale to a standard size). This preprocessing can be a significant source of errors as well as require extra implementation effort. A related advantage is that the sto-

chastic models can fit a single digit in the context of several other digits—so that precise segmentation is not required prior to classification. The primary disadvantage is that classification is slower because of the need to perform a search to fit each digit model.

Learning the Structure of Stochastic Models

All the methods I have discussed to this point learn the parameters of a stochastic model whose structure is given. An important research question is whether this structure can be learned from data as well. Recent research has made major progress in developing algorithms for this problem.

One of the first algorithms was developed by Chow and Liu (1968) for learning a network structure in the form of a directed tree. This algorithm first constructs a complete undirected graph where the nodes are the variables, and the edges are labeled with the mutual information between the variables. The algorithm then finds the maximum weighted spanning tree of this graph, chooses a root node arbitrarily, and orders the arcs to point away from the root. A nice feature of this algorithm is that it is quite fast—it runs in polynomial time.

Cooper and Herskovits (1992) developed an algorithm, called $\kappa 2$, for learning the structure of a stochastic model where all variables are observed in the training data. They adopt a maximum a posteriori (MAP) approach within a Bayesian framework. The key contribution of their paper was to derive a formula for the posterior probability of a network. This formula can be updated incrementally as changes are made to the network.

Using their formula, they implemented a simple greedy search algorithm for finding (approximately) the MAP network structure. Their algorithm requires the user to provide an ordering of the variables, and it will consider adding arcs only from variables earlier in the ordering to variables later in the ordering. This approach significantly constrains the search and also ensures that the learned network will remain acyclic. $\kappa 2$ begins with a network having no arcs—all variables are independent of one another. $\kappa 2$ evaluates the posterior probability of adding each possible single edge and makes the highest-ranking addition. This greedy algorithm is continued until no single change improves the posterior probability.

Heckerman, Geiger, and Chickering (1995) describe a modification to the $\kappa 2$ method for computing posterior probabilities and a local search algorithm that uses this improvement. Their method requires a prior probability dis-




Figure 28. Some Stages of Fitting Models to an Image of a 3 (from Revow et al. [1996]. Reproduced with permission.).

The image is displayed in the top row. The next row shows the model for a 3 being fitted. The bottom row attempts to fit the model of a 5. The light circles indicate the value of σ . In the bottom two rows, the image has been artificially thinned so that the circles are visible.

tribution in the form of a prior network and two parameters: (1) an equivalent sample size (which controls how much new data are required to override the prior) and (2) a penalty for each arc that is different from the prior network. Their local search algorithm considers all one-step changes to the network (arc addition, deletion, and reversal) and retains the change that most increases the posterior probability of the network. They obtain results comparable to $\kappa 2$. Interesting, they found that the most important role for the prior network was to provide a good starting point for local search (rather than to bias the objective function for guiding the search).

Friedman and Goldszmidt (1996) developed a network learning algorithm, called tree-augmented naive Bayes (TAN), specifically for supervised learning. The TAN algorithm starts with a naive Bayes network of the kind shown




Figure 29. Probabilistic Network Constructed by the TAN Algorithm for a Diabetes Diagnosis Task.




Figure 30. Comparison of TAN and C4.5 on 22 Benchmark Tasks.

Points above the diagonal line correspond to cases where TAN gave more accurate results than C4.5.

in figure 20 and considers adding arcs to improve the posterior probability of the network. They apply a modification of the Chow and Liu algorithm to learn a tree structure of arcs connecting the x_j variables to one another. Figure 29 shows a network learned by TAN for a diabetes diagnosis problem. Compared to figure 19, the directions of the arcs are wrong. Nonetheless, the network gives accurate classifications. Figure 30 compares the performance of TAN to C4.5 on 22 benchmark problems. The plot shows that TAN outperforms C4.5 on most of the domains.

There are many other important papers on the topic of structure learning for probabilistic networks. Four important references are (1) Verma and Pearl (1990), (2) Spirtes and Meek (1995), (3) Spiegelhalter et al. (1993), and (4)

Spirtes, Glymour, and Scheines (1993) and the references therein.

Summary: Stochastic Models

This completes my review of methods for learning with stochastic models. There are several good survey articles of this topic (Buntine 1996, 1994; Heckerman 1996).

The area of stochastic modeling is active right now. Journals and conferences that were once devoted exclusively to neural network applications are now presenting many papers on stochastic modeling. The research community is still gathering experience and developing improved algorithms for fitting and reasoning with stochastic models. Many of the stochastic models we would like to work with are intractable. The challenge is to find general-purpose, tractable approximation algorithms for reasoning with these elegant and expressive stochastic models.

Concluding Remarks

Any survey must choose particular areas and omit others. Let me briefly mention some other active areas. A central topic in machine learning is the control of overfitting. There have been many developments in this area as researchers explored various penalty functions and resampling techniques (including cross-validation) for preventing overfitting. An understanding of the overfitting process has been obtained through the statistical concepts of bias and variance, and several authors have developed bias-variance decompositions for classification problems.

Another active topic has been the study of algorithms for learning relations expressed as Horn-clause programs. This area is also known as inductive logic programming, and many algorithms and theoretical results have been developed in this area.

Finally, many papers have addressed practical problems that arise in applications such as visualization of learned knowledge, methods for extracting understandable rules from neural networks, algorithms for identifying noise and outliers in data, and algorithms for learning easy-to-understand classifiers.

There have been many exciting developments in the past five years, and the relevant literature in machine learning has been growing rapidly. As more areas within AI and computer science apply machine-learning methods to attack their problems, I expect that the flow of interesting problems and practical solutions will continue. It is an exciting time to be working in machine learning.

References

- Abu-Mostafa, Y. 1990. Learning from Hints in Neural Networks. *Journal of Complexity* 6:192–198.
- Ali, K. M., and Pazzani, M. J. 1996. Error Reduction through Learning Multiple Descriptions. *Machine Learning* 24(3): 173–202.
- Barto, A. G.; Bradtke, S. J.; and Singh, S. P. 1995. Learning to Act Using Real-Time Dynamic Programming. *Artificial Intelligence* 72:81–138.
- Barto, A. G., and Sutton, R. 1997. *Introduction to Reinforcement Learning*. Cambridge, Mass.: MIT Press.
- Bellman, R. E. 1957. *Dynamic Programming*. Princeton, N.J.: Princeton University Press.
- Bertsekas, D. P., and Tsitsiklis, J. N. 1996. *Neuro-Dynamic Programming*. Belmont, Mass.: Athena Scientific.
- Blum, A., and Rivest, R. L. 1988. Training a 3-Node Neural Network Is NP-Complete (Extended Abstract). In *Proceedings of the 1988 Workshop on Computational Learning Theory*, 9–18. San Francisco, Calif.: Morgan Kaufmann.
- Blum, A. 1997. Empirical Support for WINNOWER and Weighted-Majority Algorithms: Results on a Calendar Scheduling Domain. *Machine Learning* 26(1): 5–24.
- Breiman, L. 1996a. Bagging Predictors. *Machine Learning* 24(2): 123–140.
- Breiman, L. 1996b. Stacked Regressions. *Machine Learning* 24:49–64.
- Buntine, W. 1996. A Guide to the Literature on Learning Probabilistic Networks from Data. *IEEE Transactions on Knowledge and Data Engineering* 8:195–210.
- Buntine, W. 1994. Operations for Learning with Graphical Models. *Journal of Artificial Intelligence Research* 2:159–225.
- Buntine, W. L. 1990. A Theory of Learning Classification Rules. Ph.D. thesis, School of Computing Science, University of Technology.
- Caruana, R. 1996. Algorithms and Applications for Multitask Learning. In *Proceedings of the Thirteenth International Conference on Machine Learning*, ed. L. Saitta, 87–95. San Francisco, Calif.: Morgan Kaufmann.
- Cassandra, A. R.; Kaelbling, L. P.; and Littman, M. L. 1994. Acting Optimally in Partially Observable Stochastic Domains. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1023–1028. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Castillo, E.; Gutierrez, J. M.; and Hadi, A. 1997. *Expert Systems and Probabilistic Network Models*. New York: Springer-Verlag.
- Catlett, J. 1991. On Changing Continuous Attributes into Ordered Discrete Attributes. In *Proceedings of the European Working Session on Learning*, ed. Y. Kodratoff, 164–178. Berlin: Springer-Verlag.
- Chan, P. K., and Stolfo, S. J. 1995. Learning Arbiter and Combiner Trees from Partitioned Data for Scaling Machine Learning. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*, 39–44. Menlo Park, Calif.: AAAI Press.
- Cheeseman, P.; Self, M.; Kelly, J.; Taylor, W.; Freeman, D.; and Stutz, J. 1988. Bayesian Classification. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 607–611. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Cherkauer, K. J. 1996. Human Expert-Level Performance on a Scientific Image Analysis Task by a System Using Combined Artificial Neural Networks. In *Working Notes of the AAAI Workshop on Integrating Multiple Learned Models*, ed. P. Chan, 15–21. Available at www.cs.fit.edu/imlm/.
- Chipman, H.; George, E.; and McCulloch, R. 1996. Bayesian CART, Department of Statistics, University of Chicago. Available at gsbrem.uchicago.edu/Papers/cart.ps.
- Chow, C., and Liu, C. 1968. Approximating Discrete Probability Distributions with Dependence Trees. *IEEE Transactions on Information Theory* 14:462–467.
- Clemen, R. T. 1989. Combining Forecasts: A Review and Annotated Bibliography. *International Journal of Forecasting* 5:559–583.
- Cohen, W. W. 1995. Fast Effective Rule Induction. In *Proceedings of the Twelfth International Conference on Machine Learning*, 115–123. San Francisco, Calif.: Morgan Kaufmann.
- Cooper, G. F., and Herskovits, E. 1992. A Bayesian Method for the Induction of Probabilistic Networks from Data. *Machine Learning* 9:309–347.
- Craven, M. W., and Shavlik, J. W. 1996. Extracting Tree-Structured Representations from Trained Networks. In *Advances in Neural Information Processing Systems 8*, eds. D. S. Touretzky, M. C. Mozer, and M. Hasselmo, 24–30. Cambridge, Mass. MIT Press.
- Crites, R. H., and Barto, A. G. 1995. Improving Elevator Performance Using Reinforcement Learning. In *Advances in Neural Information Processing Systems 8*, 1017–1023. San Francisco, Calif.: Morgan Kaufmann.
- D'Ambrosio, B. 1993. Incremental Probabilistic Inference. In *Ninth Annual Conference on Uncertainty in AI*, eds. D. Heckerman and A. Mamdani, 301–308. San Francisco, Calif.: Morgan Kaufmann.
- Dayan, P., and Hinton, G. 1993. Feudal Reinforcement Learning. In *Advances in Neural Information Processing Systems 5*, 271–278. San Francisco, Calif.: Morgan Kaufmann.
- Dean, T., and Kanazawa, K. 1989. A Model for Reasoning about Persistence and Causation. *Computational Intelligence* 5(3): 142–150.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1976. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society B* 39: 1–38.
- Dietterich, T. G., and Bakiri, G. 1995. Solving Multiclass Learning Problems via Error-Correcting Output Codes. *Journal of Artificial Intelligence Research* 2:263–286.
- Dietterich, T. G., and Kong, E. B. 1995. Machine Learning Bias, Statistical Bias, and Statistical Variance of Decision Tree Algorithms, Department of Computer Science, Oregon State University. Available at ftp.cs.orst.edu/pub/tgd/papers/tr-bias.ps.gz.

- Domingos, P., and Pazzani, M. 1996. Beyond Independence: Conditions for the Optimality of the Simple Bayesian Classifier. In *Proceedings of the Thirteenth International Conference on Machine Learning*, ed. L. Saitta, 105–112. San Francisco, Calif.: Morgan Kaufmann.
- Duda, R. O., and Hart, P. E. 1973. *Pattern Classification and Scene Analysis*. New York: Wiley.
- Fayyad, U. M., and Irani, K. B. 1993. Multi-Interval Discretization of Continuous-Valued Attributes for Classification Learning. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1022–1027. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Flexner, S. B. 1983. *Random House Unabridged Dictionary*, 2d ed. New York: Random House.
- Freund, Y., and Schapire, R. E. 1995. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. In *Proceedings of the Second European Conference on Computational Learning Theory*, 23–37. Berlin: Springer-Verlag.
- Freund, Y., and Schapire, R. E. 1996. Experiments with a New Boosting Algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning*, ed. L. Saitta, 148–156. San Francisco, Calif.: Morgan Kaufmann.
- Friedman, N., and Goldszmidt, M. 1996. Building Classifiers Using Bayesian Networks. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 1277–1284. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Furnkranz, J., and Widmer, G. 1994. Incremental Reduced Error Pruning. In *Proceedings of the Eleventh International Conference on Machine Learning*, 70–77. San Francisco, Calif.: Morgan Kaufmann.
- Geman, S., and Geman, D. 1984. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6:721–741.
- Ghahramani, Z., and Jordan, M. I. 1996. Factorial Hidden Markov Models. In *Advances in Neural Information Processing Systems 8*, eds. D. S. Touretzky, M. C. Mozer, and M. Hasselmo, 472–478. Cambridge, Mass.: MIT Press.
- Gilks, W.; Thomas, A.; and Spiegelhalter, D. 1993. A Language and Program for Complex Bayesian Modelling. *The Statistician* 43:169–178.
- Golding, A. R., and Roth, D. 1996. Applying WINNOWER to Context-Sensitive Spelling Correction. In *Proceedings of the Thirteenth International Conference on Machine Learning*, ed. L. Saitta, 182–190. San Francisco, Calif.: Morgan Kaufmann.
- Gordon, G. J. 1995. Stable Function Approximation in Dynamic Programming. In *Proceedings of the Twelfth International Conference on Machine Learning*, 261–268. San Francisco, Calif.: Morgan Kaufmann.
- Hansen, L., and Salamon, P. 1990. Neural Network Ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12:993–1001.
- Hashem, S. 1993. Optimal Linear Combinations of Neural Networks. Ph.D. thesis, School of Industrial Engineering, Purdue University.
- Heckerman, D. 1996. A Tutorial on Learning with Bayesian Networks, MSR-TR-95-06, Advanced Technology Division, Microsoft Research, Redmond, Washington.
- Heckerman, D.; Geiger, D.; and Chickering, D. M. 1995. Learning Bayesian Networks: The Combination of Knowledge and Statistical Data. *Machine Learning* 20:197–243.
- Hyafil, L., and Rivest, R. L. 1976. Constructing Optimal Binary Decision Trees Is NP-Complete. *Information Processing Letters* 5(1): 15–17.
- Jensen, F. V.; Lauritzen, S. L.; and Olesen, K. G. 1990. Bayesian Updating in Recursive Graphical Models by Local Computations. *Computational Statistical Quarterly* 4:269–282.
- Jensen, F. 1996. *An Introduction to Bayesian Networks*. New York: Springer.
- John, G.; Kohavi, R.; and Pfleger, K. 1994. Irrelevant Features and the Subset Selection Problem. In *Proceedings of the Eleventh International Conference on Machine Learning*, 121–129. San Francisco, Calif.: Morgan Kaufmann.
- Jordan, M. I., and Jacobs, R. A. 1994. Hierarchical Mixtures of Experts and the EM Algorithm. *Neural Computation* 6(2): 181–214.
- Kaelbling, L. P. 1993. Hierarchical Reinforcement Learning: Preliminary Results. In *Proceedings of the Tenth International Conference on Machine Learning*, 167–173. San Francisco, Calif.: Morgan Kaufmann.
- Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. 1996. Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research* 4:237–285.
- Kanazawa, K.; Koller, D.; and Russell, S. 1995. Stochastic Simulation Algorithms for Dynamic Probabilistic Networks. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, 346–351. San Francisco, Calif.: Morgan Kaufmann.
- Kira, K., and Rendell, L. A. 1992. A Practical Approach to Feature Selection. In *Proceedings of the Ninth International Conference on Machine Learning*, 249–256. San Francisco, Calif.: Morgan Kaufmann.
- Kohavi, R., and Kunz, C. 1997. Option Decision Trees with Majority Votes. In *Proceedings of the Fourteenth International Conference on Machine Learning*, 161–169. San Francisco, Calif.: Morgan Kaufmann.
- Kohavi, R., and Sahami, M. 1996. Error-Based and Entropy-Based Discretizing of Continuous Features. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, 114–119. San Francisco, Calif.: Morgan Kaufmann.
- Kolen, J. F., and Pollack, J. B. 1991. Back Propagation Is Sensitive to Initial Conditions. In *Advances in Neural Information Processing Systems 3*, 860–867. San Francisco, Calif.: Morgan Kaufmann.
- Kong, E. B., and Dietterich, T. G. 1995. Error-Correcting Output Coding Corrects Bias and Variance. In *The Twelfth International Conference on Machine Learning*, eds. A. Prieditis and S. Russell, 313–321. San Francisco, Calif.: Morgan Kaufmann.
- Kononenko, I. 1994. Estimating Attributes: Analysis and Extensions of Relief. In *Proceedings of the 1994 European Conference on Machine Learning*, 171–182. Amsterdam: Springer Verlag.

- Kononenko, I.; Simec, E.; and Robnik-Sikonja, M. 1997. Overcoming the Myopia of Inductive Learning Algorithms with RELIEF-F. *Applied Intelligence*. Forthcoming.
- Kucera, H., and Francis, W. N. 1967. *Computational Analysis of Present-Day American English*. Providence, R.I.: Brown University Press.
- Kwok, S. W., and Carter, C. 1990. Multiple Decision Trees. In *Uncertainty in Artificial Intelligence 4*, eds. R. D. Schachter, T. S. Levitt, L. N. Kannal, and J. F. Lemmer, 327–335. Amsterdam: Elsevier Science.
- Laird, J. E.; Newell, A.; and Rosenbloom, P. S. 1987. SOAR: An Architecture for General Intelligence. *Artificial Intelligence* 33(1): 1–64.
- Littlestone, N. 1988. Learning Quickly When Irrelevant Attributes Abound: A New Linear-Threshold Algorithm. *Machine Learning* 2:285–318.
- Littman, M. L.; Cassandra, A.; and Kaelbling, L. P. 1995. Learning Policies for Partially Observable Environments: Scaling Up. In *Proceedings of the Twelfth International Conference on Machine Learning*, 362–370. San Francisco, Calif.: Morgan Kaufmann.
- Lowe, D. G. 1995. Similarity Metric Learning for a Variable-Kernel Classifier. *Neural Computation* 7(1): 72–85.
- MacKay, D. 1992. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation* 4(3): 448–472.
- Mahadevan, S. 1996. Average Reward Reinforcement Learning: Foundations, Algorithms, and Empirical Results. *Machine Learning* 22:159–195.
- Mahadevan, S., and Kaelbling, L. P. 1996. The National Science Foundation Workshop on Reinforcement Learning. *AI Magazine* 17(4): 89–97.
- McCallum, R. A. 1995. Instance-Based Utile Distinctions for Reinforcement Learning with Hidden State. In *Proceedings of the Twelfth International Conference on Machine Learning*, 387–396. San Francisco, Calif.: Morgan Kaufmann.
- Mehta, M.; Agrawal, R.; and Rissanen, J. 1996. SLIQ: A Fast Scalable Classifier for Data Mining. In *Lecture Notes in Computer Science*, 18–32. New York: Springer-Verlag.
- Merz, C. J., and Murphy, P. M. 1996. UCI Repository of Machine Learning Databases. Available at www.ics.uci.edu/mllearn/MLRepository.html.
- Miller, A. J. 1990. *Subset Selection in Regression*. New York: Chapman and Hall.
- Minton, S.; Carbonell, J. G.; Knoblock, C. A.; Kuokka, D. R.; Etzioni, O.; and Gil, Y. 1989. Explanation-Based Learning: A Problem-Solving Perspective. *Artificial Intelligence* 40:63–118.
- Moore, A. W., and Lee, M. S. 1994. Efficient Algorithms for Minimizing Cross-Validation Error. In *Proceedings of the Eleventh International Conference on Machine Learning*, eds. W. Cohen and H. Hirsh, 190–198. San Francisco, Calif.: Morgan Kaufmann.
- Munro, P., and Parmanto, B. 1997. Competition among Networks Improves Committee Performance. In *Advances in Neural Information Processing Systems 9*, 592–598. Cambridge, Mass.: MIT Press.
- Musick, R.; Catlett, J.; and Russell, S. 1992. Decision-Theoretic Subsampling for Induction on Large Databases. In *Proceedings of the Tenth International Conference on Machine Learning*, ed. P. E. Utgoff, 212–219. San Francisco, Calif.: Morgan Kaufmann.
- Neal, R. 1993. Probabilistic Inference Using Markov Chain Monte Carlo Methods, CRG-TR-93-1, Department of Computer Science, University of Toronto.
- Ok, D., and Tadepalli, P. 1996. Auto-Exploratory Average Reward Reinforcement Learning. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 881–887. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Opitz, D. W., and Shavlik, J. W. 1996. Generating Accurate and Diverse Members of a Neural-Network Ensemble. In *Advances in Neural Information Processing Systems 8*, eds. D. S. Touretzky, M. C. Mozer, and M. Hesselmo, 535–541. Cambridge, Mass.: MIT Press.
- Parmanto, B.; Munro, P. W.; and Doyle, H. R. 1996. Improving Committee Diagnosis with Resampling Techniques. In *Advances in Neural Information Processing Systems 8*, eds. D. S. Touretzky, M. C. Mozer, and M. Hesselmo, 882–888. Cambridge, Mass.: MIT Press.
- Parmanto, B.; Munro, P. W.; Doyle, H. R.; Doria, C.; Aldrighetti, L.; Marino, I. R.; Mitchel, S.; and Fung, J. J. 1994. Neural Network Classifier for Hepatoma Detection. In *Proceedings of the World Congress on Neural Networks 1*, 285–290. Hillsdale, N.J.: Lawrence Erlbaum.
- Parr, R., and Russell, S. 1995. Approximating Optimal Policies for Partially Observable Stochastic Domains. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1088–1094. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Perrone, M. P., and Cooper, L. N. 1993. When Networks Disagree: Ensemble Methods for Hybrid Neural Networks. In *Neural Networks for Speech and Image Processing*, ed. R. J. Mammone. New York: Chapman and Hall.
- Press, W. H.; Flannery, B. P.; Teukolsky, S. A.; and Vetterling, W. T. 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2d ed. Cambridge, U.K.: Cambridge University Press.
- Quinlan, J. R. 1996. Bagging, Boosting, and c4.5. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 725–730. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Quinlan, J. R. 1993. *c4.5: Programs for Empirical Learning*. San Francisco, Calif.: Morgan Kaufmann.
- Quinlan, J. R. 1990. Learning Logical Definitions from Relations. *Machine Learning* 5(3): 239–266.
- Quinlan, J. R., and Rivest, R. L. 1989. Inferring Decision Trees Using the Minimum Description Length Principle. *Information and Computation* 80:227–248.
- Rabiner, L. R. 1989. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. In *Proceedings of the IEEE* 77(2): 257–286. Washington, D.C.: IEEE Computer Society.
- Raviv, Y., and Intrator, N. 1996. Bootstrapping with Noise: An Effective Regularization Technique. *Connection Science* 8(3–4): 355–372.

- Revow, M.; Williams, C. K. I.; and Hinton, G. E. 1996. Using Generative Models for Handwritten Digit Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18(6): 592–606.
- Ricci, F., and Aha, D. W. 1997. Extending Local Learners with Error-Correcting Output Codes, Naval Center for Applied Research in Artificial Intelligence, Washington, D.C.
- Rosen, B. E. 1996. Ensemble Learning Using Decorrelated Neural Networks. *Connection Science* 8(3–4): 373–384.
- Russell, S.; Binder, J.; Koller, D.; and Kanazawa, K. 1995. Local Learning in Probabilistic Networks with Hidden Variables. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, 1146–1152. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Samuel, A. L. 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development* 3:211–229.
- Schapire, R. E. 1997. Using Output Codes to Boost Multiclass Learning Problems. In *Proceedings of the Fourteenth International Conference on Machine Learning*, 313–321. San Francisco, Calif.: Morgan Kaufmann.
- Schwartz, A. 1993. A Reinforcement Learning Method for Maximizing Undiscounted Rewards. In *Proceedings of the Tenth International Conference on Machine Learning*, ed. P. Utgoff, 298–305. San Francisco, Calif.: Morgan Kaufmann.
- Shafer, J.; Agrawal, R.; and Mehta, M. 1996. SPRINT: A Scalable Parallel Classifier for Data Mining. In *Proceedings of the Twenty-Second VLDB Conference*, 544–555. San Francisco, Calif.: Morgan Kaufmann.
- Singh, S. P. 1992. Transfer of Learning by Composing Solutions to Elemental Sequential Tasks. *Machine Learning* 8(3): 323–340.
- Singh, S., and Bertsekas, D. 1997. Reinforcement Learning for Dynamic Channel Allocation in Cellular Telephone Systems. In *Advances in Neural Information Processing Systems 10*, 974–980. Cambridge, Mass.: MIT Press.
- Smyth, P.; Heckerman, D.; and Jordan, M. I. 1997. Probabilistic Independence Networks for Hidden Markov Probability Models. *Neural Computation* 9(2): 227–270.
- Spiegelhalter, D.; Dawid, A.; Lauritzen, S.; and Cowell, R. 1993. Bayesian Analysis in Expert Systems. *Statistical Science* 8:219–282.
- Spirtes, P.; Glymour, C.; and Scheines, R. 1993. *Causation, Prediction, and Search*. New York: Springer-Verlag. Also available online at hss.cmu.edu/html/departments/philosophy/TETRAD.BOOK/book.html.
- Spirtes, P., and Meek, C. 1995. Learning Bayesian Networks with Discrete Variables from Data. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*, 294–299. San Francisco, Calif.: Morgan Kaufmann.
- Stolcke, A., and Omohundro, S. M. 1994. Best-First Model Merging for Hidden Markov Model Induction, TR-94-003, International Computer Science Institute, Berkeley, California.
- Sutton, R. S. 1988. Learning to Predict by the Methods of Temporal Differences. *Machine Learning* 3(1): 9–44.
- Tesauro, G. 1995. Temporal Difference Learning and TD-GAMMON. *Communications of the ACM* 28(3): 58–68.
- Tesauro, G. 1992. Practical Issues in Temporal Difference Learning. *Machine Learning* 8:257–278.
- Tsitsiklis, J. N., and Van Roy, B. 1996. An Analysis of Temporal-Difference Learning with Function Approximation, LIDS-P-2322, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology.
- Tumer, K., and Ghosh, J. 1996. Error Correlation and Error Reduction in Ensemble Classifiers. *Connection Science* 8(3–4): 385–404.
- Verma, T., and Pearl, J. 1990. Equivalence and Synthesis of Causal Models. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, 220–227. San Francisco, Calif.: Morgan Kaufmann.
- Watkins, C. J. C. H. 1989. Learning from Delayed Rewards. Ph.D. thesis, King's College, Oxford.
- Watkins, C. J., and Dayan, P. 1992. Technical Note: Q-LEARNING. *Machine Learning* 8:279–292.
- Wettschereck, D.; Aha, D. W.; and Mohri, T. 1997. A Review and Empirical Evaluation of Feature Weighting Methods for a Class of Lazy Learning Algorithms. *Artificial Intelligence Review* 10:1–37.
- Wettschereck, D., and Dietterich, T. G. 1995. An Experimental Comparison of the Nearest-Neighbor and Nearest-Hyperrectangle Algorithms. *Machine Learning* 19:5–27.
- Wolpert, D. 1992. Stacked Generalization. *Neural Networks* 5(2): 241–260.
- Zhang, W., and Dietterich, T. G. 1995. A Reinforcement Learning Approach to Job-Shop Scheduling. In Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, 1114–1120. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Zhang, X.; Mesirov, J. P.; and Waltz, D. L. 1992. Hybrid System for Protein Secondary Structure Prediction. *Journal of Molecular Biology* 225:1049–1063.
- Zweben, M.; Daun, B.; and Deale, M. 1994. Scheduling and Rescheduling with Iterative Repair. In *Intelligent Scheduling 8*, eds. M. Zweben and M. Fox, 241–255. San Francisco, Calif.: Morgan Kaufmann.



Thomas G. Dietterich (Ph.D., Stanford University) has been a faculty member in computer science at Oregon State University since 1985. He also held a position as senior scientist at Arris Pharmaceutical Corporation from 1991 to 1993, where he applied machine-learning methods to drug design problems. He is editor (with Jude Shavlik) of *Readings in Machine Learning* and executive editor of the journal *Machine Learning*. His e-mail address is tgd@cs.orst.edu.