




FIGURE P6.10




FIGURE P6.14

Virtues and limitations

- 6.13 (a) Compare the virtues and limitations of support vector machines with those of radial-basis function (RBF) networks with respect to the following tasks:
- (1) Pattern classification
 - (2) Nonlinear regression
- (b) Do the same for support vector machines versus multilayer perceptrons trained using the back-propagation algorithm.

Computer experiments

- 6.14 Figure P6.14 shows a set of data points corresponding to two classes, $\mathcal{C}_1$ and $\mathcal{C}_2$. Both coordinates, x_1 and x_2 , range from -1 to $+1$. Using a single radial-basis function kernel

$$K(\mathbf{x}, \mathbf{t}) = \exp(-\|\mathbf{x} - \mathbf{t}\|^2)$$

construct the optimal hyperplane and identify the support vectors for this data set.

- 6.15 The computer experiment described in Section 6.6 was for the classification of two overlapping Gaussian distributions. The following “regularization” parameter was used in that experiment: $C = 0.1$. The common width of the radial-basis functions used for constructing the inner-product kernels was $\sigma^2 = 4$. Repeat the computer experiment described in that section for the following two values of the regularization parameter:

- (a) $C = 0.05$
- (b) $C = 0.2$

Comment on your results in light of the findings reported in Section 6.6.

- 6.16 In applying radial-basis function networks to nonlinear regression problems, we often find that the use of an unlocalized basis function such as the multiquadric results in a more accurate solution than the use of a localized basis function such as the Gaussian function. It may be conjectured that a similar situation arises in the case of support vector machines, because the use of an (unbounded) polynomial learning machine may prove more accurate than a (bounded) radial-basis function machine. Using a computer experiment on a nonlinear regression problem, explore the validity of this conjecture.

Committee Machines

7.1 INTRODUCTION

In the previous three chapters we describe three different approaches to supervised learning. The MLP trained with the back-propagation algorithm, discussed in Chapter 4, relies on a form of global optimization for its design. The RBF network, discussed in Chapter 5, relies on local optimization for its design. The support vector machine, discussed in Chapter 6, exploits VC-dimension theory for its design. In this chapter we discuss another class of methods for solving supervised learning tasks. The approach used here is based on a commonly used engineering principle: divide and conquer.

According to the *principle of divide and conquer*, a complex computational task is solved by dividing it into a number of computationally simple tasks and then combining the solutions to those tasks. In supervised learning, computational simplicity is achieved by distributing the learning task among a number of *experts*, which in turn divides the input space into a set of subspaces. The combination of experts is said to constitute a *committee machine*. Basically, it fuses knowledge acquired by experts to arrive at an overall decision that is supposedly superior to that attainable by any one of them acting alone. The idea of a committee machine may be traced back to Nilsson (1965); the network structure considered therein consisted of a layer of elementary perceptrons followed by a vote-taking perceptron in the second layer.

Committee machines are universal approximators. They may be classified into two major categories:

1. *Static structures*. In this class of committee machines, the responses of several predictors (experts) are combined by means of a mechanism that does *not* involve the input signal, hence the designation “static.” This category includes the following methods:
 - *Ensemble averaging*, where the outputs of different predictors are linearly combined to produce an overall output.
 - *Boosting*, where a weak learning algorithm is converted into one that achieves arbitrarily high accuracy.
2. *Dynamic structures*. In this second class of committee machines, the input signal is directly involved in actuating the mechanism that integrates the outputs of the

individual experts into an overall output, hence the designation “dynamic.” Here we mention two kinds of dynamic structures:

- *Mixture of experts*, in which the individual responses of the experts are nonlinearly combined by means of a single gating network.
- *Hierarchical mixture of experts*, in which the individual responses of the experts are nonlinearly combined by means of several gating networks arranged in a hierarchical fashion.

In the mixture of experts, the principle of divide and conquer is applied just once, whereas in the hierarchical mixture of experts it is applied several times, resulting in a corresponding number of levels of hierarchy.

The mixture of experts and hierarchical mixture of experts may also be viewed as examples of modular networks. A formal definition of the notion of *modularity* is (Osherson et al., 1990):

A neural network is said to be modular if the computation performed by the network can be decomposed into two or more modules (subsystems) that operate on distinct inputs without communicating with each other. The outputs of the modules are mediated by an integrating unit that is not permitted to feed information back to the modules. In particular, the integrating unit both (1) decides how the outputs of the modules should be combined to form the final output of the system, and (2) decides which modules should learn which training patterns.

This definition of modularity rules out the static class of committee machines since there is no integrating unit at the output that has a decision-making role.

Organization of the Chapter

This chapter is organized in two parts. The class of static structures is covered in the first part, encompassing Sections 7.2 through 7.5. Section 7.2 discusses the method of ensemble averaging, followed by a computer experiment in Section 7.3. Section 7.4 discusses the boosting technique, followed by a computer experiment in Section 7.5.

The class of dynamic structures is covered in the second part of the chapter, encompassing Sections 7.6 through 7.13. Specifically, Section 7.6 discusses the mixture of experts (ME) as an associative Gaussian mixture model. Section 7.7 discusses the more general case, namely the hierarchical mixture of experts (HME). This latter model is closely related to standard decision trees. Then Section 7.8 describes how a standard decision tree may be used to solve the model selection problem (i.e., number of gating and expert networks) for HME. In Section 7.9, we define some *a posteriori* probabilities that assist us in the formulation of learning algorithms for the HME model. In Section 7.10 we lay the foundation for solving the parameter estimation problem by formulating the likelihood function for the HME model. Section 7.11 presents an overview of learning strategies. This is followed by a detailed discussion of the so-called EM algorithm in Section 7.12 and its application to the HME model in Section 7.13.

The chapter concludes with some final remarks in Section 7.14.

7.2 ENSEMBLE AVERAGING

Figure 7.1 shows a number of differently trained neural networks (i.e., experts), which share a common input and whose individual outputs are somehow combined to produce an overall output y . To simplify the presentation, the outputs of the experts are assumed to be scalar-valued. Such a technique is referred to as an *ensemble averaging method*¹. The motivation for its use is two-fold:

- If the combination of experts in Fig. 7.1 were replaced by a single neural network, we would have a network with a correspondingly large number of adjustable parameters. The training time for such a large network is likely to be longer than for the case of a set of experts trained in parallel.
- The risk of overfitting the data increases when the number of adjustable parameters is large compared to cardinality (i.e., size of the set) of the training data.

In any event, in using a committee machine as depicted in Fig. 7.1, the expectation is that the differently trained experts converge to different local minima on the error surface, and overall performance is improved by combining the outputs in some way.

Consider first the case of a single neural network that has been trained on a given data set. Let $\mathbf{x}$ denote an input vector not seen before, and let d denote the corresponding desired response (representing a class label or numerical response); $\mathbf{x}$ and d represent realizations of the random vector $\mathbf{X}$ and random variable D , respectively. Let $F(\mathbf{x})$ denote the input–output function realized by the network. Then, in light of the material on the bias/variance dilemma presented in Chapter 2, we may decompose the mean-square error between $F(\mathbf{x})$ and the conditional expectation $E[D|\mathbf{X} = \mathbf{x}]$ into its bias and variance components as follows:

$$E_{\mathbf{x}}[(F(\mathbf{x}) - E[D|\mathbf{X} = \mathbf{x}])^2] = B_{\mathbf{x}}(F(\mathbf{x})) + V_{\mathbf{x}}(F(\mathbf{x})) \quad (7.1)$$

where $B_{\mathbf{x}}(F(\mathbf{x}))$ is the bias squared:

$$B_{\mathbf{x}}(F(\mathbf{x})) = (E_{\mathbf{x}}[F(\mathbf{x})] - E[D|\mathbf{X} = \mathbf{x}])^2 \quad (7.2)$$

and $V_{\mathbf{x}}(F(\mathbf{x}))$ is the variance:

$$V_{\mathbf{x}}(F(\mathbf{x})) = E_{\mathbf{x}}[(F(\mathbf{x}) - E_{\mathbf{x}}[F(\mathbf{x})])^2] \quad (7.3)$$




FIGURE 7.1 Block diagram of a committee machine based on ensemble averaging.

The expectation $E_{\mathcal{D}}$ is taken over the space $\mathcal{D}$, defined as the *space encompassing the distribution of all training sets (i.e., inputs and target outputs) and the distribution of all initial conditions*.

There are different ways of individually training the expert networks in Fig. 7.1 and also different ways of combining their outputs. For the discussion presented here, we will consider the situation where the expert networks have an identical configuration, but they are trained starting from different initial conditions. For the combiner at the output of the committee machine in Fig. 7.1, we will use a simple *ensemble average*.² Let $\mathcal{J}$ denote the *space of all initial conditions*. Let $F_I(\mathbf{x})$ denote the average of the input-output functions of the expert networks in Fig. 7.1 over a "representative" number of initial conditions. By analogy with Eq. (7.1), we may write

$$E_{\mathcal{J}}[(F_I(\mathbf{x}) - E[D | \mathbf{X} = \mathbf{x}])^2] = B_{\mathcal{J}}(F(\mathbf{x})) + V_{\mathcal{J}}(F(\mathbf{x})) \quad (7.4)$$

where $B_{\mathcal{J}}(F(\mathbf{x}))$ is the squared bias defined over the space $\mathcal{J}$:

$$B_{\mathcal{J}}(F(\mathbf{x})) = (E_{\mathcal{J}}[F_I(\mathbf{x})] - E[D | \mathbf{X} = \mathbf{x}])^2 \quad (7.5)$$

and $V_{\mathcal{J}}(F(\mathbf{x}))$ is the corresponding variance:

$$V_{\mathcal{J}}(F(\mathbf{x})) = E_{\mathcal{J}}[(F_I(\mathbf{x}) - E_{\mathcal{J}}[F_I(\mathbf{x})])^2] \quad (7.6)$$

The expectation $E_{\mathcal{J}}$ is taken over the space $\mathcal{J}$.

From the definition of space $\mathcal{D}$, we may view it as the product of the space of initial conditions, $\mathcal{J}$, and the *remnant space* denoted by $\mathcal{D}'$. Accordingly, we may write, again by analogy with Eq. (7.1):

$$E_{\mathcal{D}'}[(F_I(\mathbf{x}) - E[D | \mathbf{X} = \mathbf{x}])^2] = B_{\mathcal{D}'}(F_I(\mathbf{x})) + V_{\mathcal{D}'}(F_I(\mathbf{x})) \quad (7.7)$$

where $B_{\mathcal{D}'}(F_I(\mathbf{x}))$ is the squared bias defined over the remnant space $\mathcal{D}'$:

$$B_{\mathcal{D}'}(F_I(\mathbf{x})) = (E_{\mathcal{D}'}[F_I(\mathbf{x})] - E[D | \mathbf{X} = \mathbf{x}])^2 \quad (7.8)$$

and $V_{\mathcal{D}'}(F_I(\mathbf{x}))$ is the corresponding variance:

$$V_{\mathcal{D}'}(F_I(\mathbf{x})) = E_{\mathcal{D}'}[(F_I(\mathbf{x}) - E_{\mathcal{D}'}[F_I(\mathbf{x})])^2] \quad (7.9)$$

From the definitions of spaces $\mathcal{D}$, $\mathcal{J}$, and $\mathcal{D}'$, we readily see that

$$E_{\mathcal{D}}[F_I(\mathbf{x})] = E_{\mathcal{D}}[F(\mathbf{x})] \quad (7.10)$$

It follows therefore that Eq. (7.8) may be rewritten in the equivalent form:

$$\begin{aligned} B_{\mathcal{D}'}(F_I(\mathbf{x})) &= (E_{\mathcal{D}}[F(\mathbf{x})] - E[D | \mathbf{X} = \mathbf{x}])^2 \\ &= B_{\mathcal{D}}(F(\mathbf{x})) \end{aligned} \quad (7.11)$$

Consider next the variance $V_{\mathcal{D}'}(F_I(\mathbf{x}))$ of Eq. (7.9). Since the variance of a random variable is equal to the mean-square value of that random variable minus its bias squared, we may equivalently write

$$\begin{aligned} V_{\mathcal{D}'}(F_I(\mathbf{x})) &= E_{\mathcal{D}'}[(F_I(\mathbf{x}))^2] - (E_{\mathcal{D}'}[F_I(\mathbf{x})])^2 \\ &= E_{\mathcal{D}'}[(F_I(\mathbf{x}))^2] - (E_{\mathcal{D}}[F(\mathbf{x})])^2 \end{aligned} \quad (7.12)$$

where in the last line we have made use of Eq. (7.10). Similarly, we may redefine Eq. (7.3) in the equivalent form

$$V_{\mathcal{D}}(F_I(\mathbf{x})) = E_{\mathcal{D}}[(F(\mathbf{x}))^2] - (E_{\mathcal{D}}[F(\mathbf{x})])^2 \quad (7.13)$$

Note that the mean-square value of the function $F(\mathbf{x})$ over the entire space $\mathcal{D}$ is destined to be equal to or greater than the mean-square value of the ensemble-averaged function $F_I(\mathbf{x})$ over the remnant space $\mathcal{D}'$. That is,

$$E_{\mathcal{D}}[F(\mathbf{x})^2] \geq E_{\mathcal{D}'}[(F_I(\mathbf{x}))^2]$$

In light of this inequality, comparison of Eqs. (7.12) and (7.13) immediately reveals that

$$V_{\mathcal{D}'}(F_I(\mathbf{x})) \leq V_{\mathcal{D}}(F(\mathbf{x})) \quad (7.14)$$

Thus, from Eqs. (7.11) and (7.14) we draw two conclusions (Naftaly et al., 1997):

1. The bias of the ensemble-averaged function $F_I(\mathbf{x})$, pertaining to the committee machine of Fig. 7.1, is exactly the same as that of the function $F(\mathbf{x})$ pertaining to a single neural network.
2. The variance of the ensemble-averaged function $F_I(\mathbf{x})$ is less than that of the function $F(\mathbf{x})$.

These theoretical findings point to a training strategy for reducing the overall error produced by a committee machine due to *varying initial conditions* (Naftaly et al., 1997): The constituent experts of the machine are purposely *overtrained*, the use of which is justified on the following grounds. Insofar as the individual experts are concerned, the bias is reduced at the cost of variance. Subsequently, however, the variance is reduced by ensemble averaging the experts over the initial conditions, leaving the bias unchanged.

7.3 COMPUTER EXPERIMENT I

In this computer experiment on the ensemble-averaging method, we revisit the pattern classification problem considered in the previous three chapters. The problem pertains to the classification of two overlapping two-dimensional Gaussian distributions. The two distributions have different mean vectors and different variances. The statistics of distribution 1 (class $\mathcal{C}_1$) are

$$\mu_1 = [0, 0]^T$$

$$\sigma_1^2 = 1$$

The statistics of distribution 2 (class $\mathcal{C}_2$) are

$$\mu_2 = [2, 0]^T$$

$$\sigma_2^2 = 4$$

Scatter plots for these two distributions are shown in Fig. 4.13.

The two classes are assumed to be equiprobable. The costs for misclassifications are assumed to be equal, and the costs for correct classifications are assumed to be zero. On this basis, the (optimum) Bayes classifier achieves a probability of correct

classification $p_c = 81.51$ percent. Details of this calculation are also presented in Chapter 4.

In the computer experiment described in Chapter 4 we were able to achieve a probability of correct classification close to 80 percent using a multilayer perceptron with two hidden neurons and trained using the back-propagation algorithm. In this experiment we study a committee machine composed as follows:

- Ten experts.
- Each expert made up of a multilayer perceptron with two hidden neurons.

All the experts were individually trained using the back-propagation algorithm. The parameters used in the algorithm were

Learning-rate parameter, $\eta = 0.1$

Momentum constant, $\alpha = 0.5$

The size of the training sample was 500 patterns. All the experts were trained on the same data set, but were initialized differently. In particular, the initial values of the synaptic weights and thresholds were picked at random from a uniform distribution inside the range $[-1, 1]$.

Table 7.1 presents a summary of the classification performances of the 10 experts trained on 500 patterns using the test set. The probability of correct classification obtained by simply taking the arithmetic average of the 10 results presented in Table 7.1 is $p_{c,av} = 79.37$ percent. On the other hand, by using the ensemble-averaging method, that is, by simply summing the individual outputs of the 10 experts and then computing the probability of correct classification, we obtained the result: $p_{c,ens} = 80.27$ percent. This result represents an improvement of 0.9 percent over $p_{c,av}$. The improvement of $p_{c,ens}$ over $p_{c,av}$ was maintained in all the trials of the experiment. The classification results were all computed using 32,000 test patterns.

TABLE 7.1 Classification Performances of Individual Experts Used in a Committee Machine

Expert	Correct classification percentage
Net 1	80.65
Net 2	76.91
Net 3	80.06
Net 4	80.47
Net 5	80.44
Net 6	76.89
Net 7	80.55
Net 8	80.47
Net 9	76.91
Net 10	80.38

Summarizing the results of this experiment, we may say: The classification performance is improved by overtraining the individual multilayer perceptrons (experts), summing their individual numerical outputs to produce the overall output of the committee machine, and then making a decision.

7.4 BOOSTING

As mentioned in the introduction, boosting is another method that belongs to the “static” class of committee machines. Boosting is quite different from ensemble averaging. In a committee machine based on ensemble averaging, all the experts in the machine are trained on the same data set; they may differ from each other in the choice of initial conditions used in network training. By contrast, in a boosting machine the experts are trained on data sets with entirely different distributions; it is a general method that can be used to improve the performance of *any* learning algorithm.

*Boosting*³ can be implemented in three fundamentally different ways:

1. *Boosting by filtering.* This approach involves filtering the training examples by different versions of a weak learning algorithm. It assumes the availability of a large (in theory, infinite) source of examples, with the examples being either discarded or kept during training. An advantage of this approach is that it allows for a small memory requirement compared to the other two approaches.
2. *Boosting by subsampling.* This second approach works with a training sample of *fixed* size. The examples are “resampled” according to a given probability distribution during training. The error is calculated with respect to the fixed training sample.
3. *Boosting by reweighting.* This third approach also works with a fixed training sample, but it assumes that the weak learning algorithm can receive “weighted” examples. The error is calculated with respect to the weighted examples.

In this section we describe two different boosting algorithms. One algorithm, due to Schapire (1990), belongs to approach 1. The other algorithm, known as AdaBoost due to Freund and Schapire (1996a, 1996b), belongs to approach 2.

Boosting by Filtering

The original idea of boosting described in Schapire (1990) is rooted in a *distribution-free or probably approximately correct (PAC) model of learning*. From the discussion of PAC learning presented in Chapter 2, we recall that a *concept* is a Boolean function in some domain of instances that contains encodings of all objects of interest. In PAC learning, a learning machine tries to identify an unknown binary concept on the basis of randomly chosen examples of the concept. To be more specific, the goal of the learning machine is to find a hypothesis or prediction rule with an error rate of at most ϵ , for arbitrarily small positive values of ϵ , and this should hold uniformly for all input distributions. It is for this reason that the PAC learning model is also referred to as a *strong learning model*. Since the examples are of a random nature, it is likely that the learning

machine will be unable to learn anything about the unknown concept due to the presentation of a highly unrepresentative example. We therefore require the learning model to succeed only in finding a good approximation to the unknown concept with a probability of $1 - \delta$, where δ is a small positive number.

In a variant of the PAC learning model, called a *weak learning model*, the requirement to learn an unknown concept is relaxed dramatically. The learning machine is now required to find a hypothesis with an error rate only slightly less than $1/2$. When a hypothesis guesses a binary label in an entirely random manner on every example, it can be right or wrong with equal probability. That is, it achieves an error rate of exactly $1/2$. It follows therefore that a weak learning model has to perform only slightly better than random guessing. The notion of weak learnability was introduced by Kearns and Valiant (1989), who posed the *hypothesis boosting problem* that is embodied in the following question:

Are the notions of strong and weak learning equivalent?

In other words, is any concept class that is weakly learnable also strongly learnable? This question, which is perhaps surprising, was answered in the affirmative by Schapire (1990). The proof presented therein was constructive. Specifically, an algorithm was devised for directly converting a weak learning model into a strong learning model. This is achieved by modifying the distribution of examples in such a way that a strong learning model is built around a weak one.

In boosting by filtering, the committee machine consists of three experts or sub-hypotheses. The algorithm used to train them is called a *boosting algorithm*. The three experts are arbitrarily labeled “first,” “second,” and “third.” The three experts are individually trained as follows:

1. The first expert is trained on a set consisting of N_1 examples.
2. The trained first expert is used to *filter* another set of examples by proceeding in the following manner:
 - Flip a fair coin; this in effect simulates a random guess.
 - If the result is *heads*, pass new patterns through the first expert and discard correctly classified patterns until a pattern is misclassified. That misclassified pattern is added to the training set for the second expert.
 - If the result is *tails*, do the opposite. Specifically, pass new patterns through the first expert and discard incorrectly classified patterns until a pattern is classified correctly. That correctly classified pattern is added to the training set for the second expert.
 - Continue this process until a total of N_1 examples has been filtered by the first expert. This set of filtered examples constitutes the training set for the second expert.

By following this coin flipping procedure, it is ensured that if the first expert is tested on the second set of examples, it would have an error rate of $1/2$. In other words, the second set of N_1 examples available for training the second expert has a distribution entirely different from the first set of N_1 examples used to train the first expert. In this way, the second expert is forced to learn a distribution different from that learned by the first expert.

3. Once the second expert has been trained in the usual way, a third training set is formed for the third expert by proceeding in the following manner:
 - Pass a new pattern through both the first and second experts. If the two experts agree in their decisions, discard that pattern. If, on the other hand, they disagree, the pattern is added to the training set for the third expert.
 - Continue with this process until a total of N_1 examples has been filtered jointly by the first and second experts. This set of jointly filtered examples constitutes the training set for the third expert.

The third expert is then trained in the usual way, and the training of the entire committee machine is thereby completed.

This three-point filtering procedure is illustrated in Fig. 7.2.

Let N_2 denote the number of examples that must be filtered by the first expert to obtain the training set of N_1 examples for the second expert. Note that N_1 is fixed, and N_2 depends on the generalization error rate of the first expert. Let N_3 denote the number of examples that must be jointly filtered by the first and second experts to obtain the training set of N_1 examples for the third expert. With N_1 examples also needed to train the first expert, the total size of data set needed to train the entire committee machine is $N_4 = N_1 + N_2 + N_3$. However, the computational cost is based on $3N_1$ examples because N_1 is the number of examples actually used to train each of the three experts. We may therefore say that the boosting algorithm described herein is indeed “smart” in the sense that the committee machine requires a large set of examples for its operation, but only a subset of that data set is used to perform the actual training.

Another noteworthy point is that the filtering operation performed by the first expert and the joint filtering operation performed by the first and second experts make the second and third experts, respectively, focus on “hard-to-learn” parts of the distribution.




FIGURE 7.2 Illustration of boosting by filtering.

In the theoretical derivation of the boosting algorithm originally presented in Schapire (1990), simple *voting* was used to evaluate the performance of the committee machine on test patterns not seen before. Specifically, a test pattern is presented to the committee machine. If the first and second experts in the committee machine agree in their respective decisions, that class label is used. Otherwise, the class label discovered by the third expert is used. However, in experimental work presented in Drucker et al. (1993, 1994), it has been determined that *addition* of the respective outputs of the three experts yields a better performance than voting. For example, in the optical character recognition (OCR) problem, the addition operation is performed simply by adding the "digit 0" outputs of the three experts, and likewise for the other nine digit outputs.

Suppose that the three experts (i.e., subhypotheses) have an error rate of $\epsilon < 1/2$ with respect to the distributions on which they are individually trained; that is, all three of them are weak learning models. In Schapire (1990), it is proved that the overall error rate of the committee machine is *bounded* by

$$g(\epsilon) = 3\epsilon^2 - 2\epsilon^3 \quad (7.15)$$

The bound $g(\epsilon)$ is plotted versus ϵ in Fig. 7.3. From this figure we see that the bound is significantly smaller than the original error rate ϵ . By applying the boosting algorithm recursively, the error rate can be made arbitrarily small. In other words, a weak learning model, which performs only slightly better than random guessing, is converted into a strong learning model. It is in that sense that we may say that strong and weak learnability are indeed equivalent.

AdaBoost

A practical limitation of boosting by filtering is that it often requires a large training sample. This limitation can be overcome by using another boosting algorithm called *AdaBoost* (Freund and Schapire, 1996a, 1996b), which belongs to boosting by resampling. The sampling framework of AdaBoost is the natural framework of batch learning; most importantly, it permits the training data to be reused.




FIGURE 7.3 Graph of Eq. (7.15) for boosting by filtering.

As with the boosting by filtering algorithm, AdaBoost has access to a weak learning model. The goal of the new algorithm is to find a final mapping function or hypothesis with low error rate relative to a given distribution $\mathcal{D}$ over the labeled training examples. It differs from other boosting algorithms in two respects:

- AdaBoost adjusts *adaptively* to the errors of the weak hypothesis returned by the weak learning model, hence the name of the algorithm.
- The bound on performance of AdaBoost depends only on the performance of the weak learning model on those distributions that are actually generated during the learning process.

AdaBoost operates as follows. On iteration n , the boosting algorithm provides the weak learning model with a distribution $\mathcal{D}_n$ over the training sample $\mathcal{T}$. In response, the weak learning model computes a hypothesis $\mathcal{F}_n : \mathbf{X} \rightarrow Y$ that correctly classifies a fraction of the training examples. The error is measured with respect to the distribution $\mathcal{D}_n$. The process continues for T iterations, and finally the boosting machine combines the hypotheses $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_T$ into a single final hypothesis $\mathcal{F}_{\text{fin}}$.

To calculate (1) the distribution $\mathcal{D}_n$ on iteration n , and (2) the final hypothesis $\mathcal{F}_{\text{fin}}$, the simple procedure summarized in Table 7.2 is used. The initial distribution $\mathcal{D}_1$ is uniform over the training sample $\mathcal{T}$, as shown by

$$\mathcal{D}_1(i) = \frac{1}{n} \quad \text{for all } i$$

Given the distribution $\mathcal{D}_n$ and weak hypothesis $\mathcal{F}_n$ on iteration n of the algorithm, the next distribution $\mathcal{D}_{n+1}$ is computed by multiplying the weight of example i by some number $\beta_n \in [0, 1)$ if $\mathcal{F}_n$ classifies the input vector $\mathbf{x}_i$ correctly; otherwise, the weight is left unchanged. The weights are then renormalized by dividing by the normalizing constant Z_n . In effect, the "easy" examples in the training set $\mathcal{T}$ that are correctly classified by many of the previous weak hypotheses are given lower weights, whereas the "hard" examples that are often misclassified are given higher weights. Thus the AdaBoost algorithm focuses the most weight on those examples that appear to be hardest for it to classify.

As for the final hypothesis $\mathcal{F}_{\text{fin}}$, it is computed as a weighted vote (i.e., weighted linear threshold) of the weak hypotheses $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_T$. That is, for a given input vector $\mathbf{x}$, the final hypothesis $\mathcal{F}_{\text{fin}}$ outputs the label d that maximizes the sum of the weights of the weak hypotheses predicting that label. The weight of hypothesis $\mathcal{F}_n$ is defined to be $\log(1/\beta_n)$, so that greater weight is given to hypotheses with lower error.

The important theoretical property of AdaBoost is stated in the following theorem (Freund and Schapire, 1996a):

Suppose the weak learning model, when called by AdaBoost, generates hypotheses with errors $\epsilon_1, \epsilon_2, \dots, \epsilon_T$, where the error ϵ_n on iteration n of the AdaBoost algorithm is defined by

$$\epsilon_n = \sum_{i: \mathcal{F}_n(\mathbf{x}_i) \neq d_i} \mathcal{D}_n(i)$$

Assume that $\epsilon_n \leq 1/2$, and let $\gamma_n = 1/2 - \epsilon_n$. Then the following upper bound holds on the error of the final hypothesis:

$$\frac{1}{N} |\{i: \mathcal{F}_{\text{fin}}(\mathbf{x}_i) \neq d_i\}| \leq \prod_{n=1}^T \sqrt{1 - 4\gamma_n^2} \leq \exp\left(-2 \sum_{n=1}^T \gamma_n^2\right) \quad (7.16)$$

TABLE 7.2 Summary of AdaBoost

Input: Training sample $\{(x_n, d_n)\}_{n=1}^N$
 Distribution $\mathcal{G}$ over the N labeled examples
 Weak learning model
 Integer T specifying the number of iterations of the algorithm

Initialization: Set $\mathcal{G}_1(i) = 1/N$ for all i

Computation: Do the following for $n = 1, 2, \dots, T$:

1. Call the weak learning model, providing it with the distribution $\mathcal{G}_n$.
2. Get back hypothesis $\mathcal{F}_n: \mathbf{X} \rightarrow Y$.
3. Calculate the error of hypothesis $\mathcal{F}_n$:

$$\epsilon_n = \sum_{i: \mathcal{F}_n(\mathbf{x}_i) \neq d_i} \mathcal{G}_n(i)$$

4. Set $\beta_n = \frac{\epsilon_n}{1 - \epsilon_n}$.

5. Update the distribution $\mathcal{G}_n$:

$$\mathcal{G}_{n+1}(i) = \frac{\mathcal{G}_n(i)}{Z_n} \times \begin{cases} \beta_n & \text{if } \mathcal{F}_n(\mathbf{x}_i) = d_i \\ 1 & \text{otherwise} \end{cases}$$

where Z_n is a normalization constant (chosen so that $\mathcal{G}_{n+1}(i)$ is a probability distribution).

Output: The final hypothesis is

$$\mathcal{F}_n(\mathbf{x}) = \arg \max_{d \in Y} \sum_{n: \mathcal{F}_n(\mathbf{x}) = d} \log \frac{1}{\beta_n}$$

This theorem shows that if the weak hypotheses constructed by the weak learning model consistently have error only slightly better than $1/2$, then the training error of the final hypothesis $\mathcal{F}_{\text{fin}}$ drops to zero exponentially fast. However, this does not mean that the generalization error on test data is necessarily small. Experiments presented in Freund and Schapire (1996a) indicate two things. First, the theoretical bound on the training error is often weak. Second, the generalization error tends to be much better than what the theory would suggest.

Table 7.2 presents a summary of AdaBoost for a binary classification problem.

When the number of possible classes (labels) is $M > 2$, the boosting problem becomes more intricate because the probability that random guessing gives the correct label is $1/M$, which is now smaller than $1/2$. For boosting to be able to use any hypothesis that is slightly better than random guessing in such a situation, we need to somehow change the algorithm and the definition of what is meant to be a “weakly learning” algorithm. Ways of invoking that change are described in Freund and Schapire (1997) and Schapire (1997).

Error Performance

Experiments with AdaBoost reported in Breiman (1996b) show that when the training error and test error are plotted as a function of the number of boosting iterations, we often find that the test error continues to decrease after the training error has reduced




FIGURE 7.4 Conceptualized error performance of the AdaBoost algorithm.

essentially to zero. This phenomenon is illustrated in Fig. 7.4. A similar result was reported earlier in Drucker et al. (1994) for boosting by filtering.

The phenomenon displayed in Fig. 7.4 is very surprising in light of what we know about the generalization performance of a single neural network. From Chapter 4 we recall that in the case of a multilayer perceptron trained with the back-propagation algorithm, the error on test (validation) data decreases, reaches a minimum, and then increases due to overfitting; see Fig. 4.20. The behavior displayed in Fig. 7.4 is remarkably different in that as the networks become more and more complex through increased training, the generalization error continues to decrease. Such a phenomenon appears to contradict *Occam's razor*, which states that a learning machine should be as simple as possible in order to achieve a good generalization performance.

In Schapire et al. (1997), an explanation is given for this phenomenon as it applies to AdaBoost. The key idea of the analysis presented therein is that in evaluating the generalization error produced by a boosting machine, not only the training error should be considered but also the *confidence* of classifications. The analysis presented therein reveals a relation between boosting and support vector machines; support vector machines are considered in the previous chapter. In particular, the classification *margin*, for example, is defined as the difference between the weight assigned to the correct label pertaining to that example and the maximal weight assigned to any single incorrect label. From this definition, it is easy to see that the margin is a number in the range $[-1, 1]$ and that an example is correctly classified if and only if its margin is positive. Thus Schapire et al. show that the phenomenon observed in Fig. 7.4 is indeed related to the distribution of margins of the training examples with respect to the generated voting classification error. It should be emphasized again that

the margin analysis presented in Schapire et al. (1997) is specific to AdaBoost and does not apply to other boosting algorithms.

COMPUTER EXPERIMENT II

In this experiment we explore the boosting by filtering algorithm to solve a fairly difficult pattern classification task. The classification problem is two-dimensional, involving nonconvex decision regions, as shown in Fig. 7.5. One class of patterns consists of data points lying inside the region labeled $\mathcal{C}_1$, and the other class of patterns consists of data points inside the region labeled $\mathcal{C}_2$. The requirement is to design a committee machine that decides whether a test pattern belongs to class $\mathcal{C}_1$ or to class $\mathcal{C}_2$.

The committee machine used to solve this problem consists of three experts. Each expert consists of a 2-5-2 multilayer perceptron that has two input nodes, five hidden neurons, and two output neurons. The back-propagation algorithm was used to perform the training. Figure 7.6 displays the scatter plots of the data used to train the three experts. The data shown in Fig. 7.6a were used to train expert 1. The data shown in Fig. 7.6b were filtered by expert 1 after its training was completed; this data set was used to train expert 2. The data shown in Fig. 7.6c were filtered by the combined action of trained experts 1 and 2; this data set was used to train expert 3. The size of training sample for each expert was $N_1 = 1000$ patterns. Examining these three figures we observe:

- The training data for expert 1 in Fig. 7.6a are uniformly distributed.
- The training data for expert 2 in Fig. 7.6b exhibit concentrations of data points in areas labeled A and B that are seemingly difficult for the first expert to classify. The number of data points in these two regions is equal to the number of the correctly classified points.




FIGURE 7.5 Pattern configurations for experiment on boosting.




FIGURE 7.6 Scatter plots for expert training in computer experiment on boosting: (a) Expert 1. (b) Expert 2. (c) Expert 3.

- The training data for expert 3 in Fig. 7.6c exhibit an even greater concentration of data points seemingly difficult for both experts 1 and 2 to classify.

Figures 7.7a, 7.7b, and 7.7c display the decision boundaries formed by experts 1, 2, and 3, respectively. Figure 7.7d displays the overall decision boundary formed by the combined action of all three experts, which is obtained by simply summing their individual outputs. Note that the difference between the decision regions of Figs. 7.7a and 7.7b pertaining to experts 1 and 2 defines the distribution of data points in Fig. 7.6c used to train expert 3.

The probabilities of correct classification for the three experts on test data were:

Expert 1 : 75.15 percent
 Expert 2 : 71.44 percent
 Expert 3 : 68.90 percent




FIGURE 7.7 Decision boundaries formed by the different experts in the boosting experiment. (a) Expert 1. (b) Expert 2. (c) Expert 3. (d) Entire committee machine.

The overall probability of correct classification for the entire committee machine was 91.79 percent, which was computed using 32,000 patterns for test data. The overall decision boundary constructed by the boosting algorithm for the three experts shown in Fig. 7.7d is further evidence for this good classification performance.

7.6 ASSOCIATIVE GAUSSIAN MIXTURE MODEL

In the second part of the chapter, beginning with this section, we study the second class of committee machines, namely dynamic structures. The term “dynamic” is used here in the sense that integration of knowledge acquired by the experts is accomplished under the action of the input signal.

To begin the discussion, consider a modular network in which the learning process proceeds by fusing self-organized and supervised forms of learning in a seam-

less fashion. The experts are technically performing supervised learning in that their individual outputs are combined to model the desired response. There is, however, a sense in which the experts are also performing self-organized learning; that is, they self-organize to find a good partitioning of the input space so that each expert does well at modeling its own subspace, and as a whole group they model the input space well.

In the learning scheme just described, there is a point of departure from the schemes considered in the previous three chapters in that a specific model is assumed for generating the training data.

Probabilistic Generative Model

To fix ideas, consider a *regression* problem in which a regressor $\mathbf{x}$ produces a response denoted by random variable D ; a realization of this random variable is denoted by d . Without loss of generality, we have adopted a scalar form of regression, merely to simplify the presentation. Specifically, we assume that the generation of response d is governed by the following probabilistic model (Jordan and Jacobs, 1995):

1. An input vector $\mathbf{x}$ is picked at random from some prior distribution.
2. A particular rule, say the k th rule, is selected in accordance with the conditional probability $P(k|\mathbf{x}, \mathbf{a}^{(0)})$, given $\mathbf{x}$ and some parameter vector $\mathbf{a}^{(0)}$.
3. For rule k , $k = 1, 2, \dots, K$, the model response d is linear in $\mathbf{x}$, with an additive error ϵ_k modeled as a Gaussian distributed random variable with zero mean and unit variance:

$$E[\epsilon_k] = 0 \quad \text{for all } k \quad (7.17)$$

and

$$\text{var}[\epsilon_k] = 1 \quad \text{for all } k \quad (7.18)$$

Under point 3, the unity variance assumption is made just for didactic simplicity. In general, each expert can have a different output variance that can be learned from the training data.

The probabilistic generation of D is determined by the conditional probability $P(D = d|\mathbf{x}, \mathbf{w}_k^{(0)})$, given $\mathbf{x}$ and some parameter vector $\mathbf{w}_k^{(0)}$, for $k = 1, 2, \dots, K$. We do not require that the probabilistic generative model just described must have a direct correspondence to a physical reality. Rather, we simply require that the probabilistic decisions embodied therein represent an *abstract* model, which with increasing precision specifies the location of the *conditional mean of response d on a nonlinear manifold* that relates the input vector to mean output (Jordan, 1994).

According to this model, the response D can be generated in K different ways, corresponding to the K choices of label k . Thus, the conditional probability of generating response $D = d$, given input vector $\mathbf{x}$, is equal to

$$P(D = d|\mathbf{x}, \mathbf{0}^{(0)}) = \sum_{k=1}^K P(D = d|\mathbf{x}, \mathbf{w}_k^{(0)}) P(k|\mathbf{x}, \mathbf{a}^{(0)}) \quad (7.19)$$

where $\mathbf{0}^{(0)}$ is the *generative model parameter vector* denoting the combination of $\mathbf{a}^{(0)}$ and $\{\mathbf{w}_k^{(0)}\}_{k=1}^K$. The superscript 0 in $\mathbf{a}^{(0)}$ and $\mathbf{w}_k^{(0)}$ is intended to distinguish the generative model parameters from those of the mixture of experts model considered next.

Mixture of Experts Model

Consider the network configuration of Fig. 7.8, referred to as a *mixture of experts (ME) model*.⁴ Specifically, it consists of K supervised modules called *expert networks* or simply *experts*, and an integrating unit called a *gating network* that performs the function of a mediator among the expert networks. It is assumed here that the different experts work best in different regions of the input space in accordance with the probabilistic generative model described, hence the need for the gating network.

With the regression problem assumed to be scalar, each expert network consists of a linear filter. Figure 7.9 shows the signal-flow graph of a single neuron constituting expert k . Thus, the output produced by expert k is the inner product of the input vector $\mathbf{x}$ and synaptic weight vector $\mathbf{w}_k$ of this neuron, as shown by

$$y_k = \mathbf{w}_k^T \mathbf{x}, \quad k = 1, 2, \dots, K \quad (7.20)$$




FIGURE 7.8 Block diagram of the ME model; the scalar outputs of the experts are mediated by a gating network.




FIGURE 7.9 Signal-flow graph of a single linear neuron constituting expert k .




FIGURE 7.10 (a) Single layer of softmax neurons for the gating network. (b) Signal-flow graph of a softmax neuron.

The gating network consists of a single layer of K neurons, with each neuron assigned to a specific expert. Figure 7.10a shows the architectural graph of the gating network, and Fig. 7.10b shows the signal-flow graph of neuron k in that network. Unlike the experts, the neurons of the gating network are nonlinear, with their activation functions defined by

$$g_k = \frac{\exp(u_k)}{\sum_{j=1}^K \exp(u_j)}, \quad k = 1, 2, \dots, K \quad (7.21)$$

where u_k is the inner product of the input vector $\mathbf{x}$ and synaptic weight vector $\mathbf{a}_k$; that is,

$$u_k = \mathbf{a}_k^T \mathbf{x}, \quad k = 1, 2, \dots, K \quad (7.22)$$

The “normalized” exponential transformation of Eq. (7.21) may be viewed as a multi-input generalization of the logistic function. It preserves the rank order of its input values, and is a differentiable generalization of the “winner-takes-all” operation of picking the maximum value. For this reason, the activation function of Eq. (7.21) is referred to as *softmax* (Bridle, 1990a). Note that the linear dependence of u_k on the input $\mathbf{x}$ makes the outputs of the gating network nonlinear functions of $\mathbf{x}$.

For a probabilistic interpretation of the gating network’s role, we may view it as a “classifier” that maps the input vector $\mathbf{x}$ into *multinomial probabilities* so that the different experts will be able to match the desired response (Jordan and Jacobs, 1995).

Most importantly, the use of softmax as the activation function for the gating network ensures that these probabilities satisfy the following requirements:

$$0 \leq g_k \leq 1 \quad \text{for all } k \quad (7.23)$$

and

$$\sum_{k=1}^K g_k = 1 \quad (7.24)$$

Let y_k denote the output of the k th expert in response to the input vector $\mathbf{x}$. The overall output of the ME model is

$$y = \sum_{k=1}^K g_k y_k \quad (7.25)$$

where, as pointed out earlier, g_k is a nonlinear function of $\mathbf{x}$. Given that rule k of the probabilistic model is selected and the input is $\mathbf{x}$, an individual output y_k is treated as the conditional mean of the random variable D , as shown by

$$\begin{aligned} E[D | \mathbf{x}, k] &= y_k \\ &= \mathbf{w}_k^T \mathbf{x}, \quad k = 1, 2, \dots, K \end{aligned} \quad (7.26)$$

With μ_k denoting the conditional mean of D , we may write

$$\mu_k = y_k, \quad k = 1, 2, \dots, K \quad (7.27)$$

The variance of D is the same as that of the error ϵ_k . Thus, invoking the use of Eq. (7.18), we may write

$$\text{var}[D | \mathbf{x}, k] = 1, \quad k = 1, 2, \dots, K \quad (7.28)$$

The probability density function of D , given the input vector $\mathbf{x}$ and given that the k th rule of the probabilistic generative model (i.e., expert k) is selected, may therefore be described as:

$$f_D(d | \mathbf{x}, k, \mathbf{0}) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(d - y_k)^2\right), \quad k = 1, 2, \dots, K \quad (7.29)$$

where $\mathbf{0}$ is a parameter vector denoting the parameters of the gating network and those of the experts in the ME model. The probability density function of D , given $\mathbf{x}$, is the *mixture* of the probability density functions $\{f_D(d | \mathbf{x}, k, \mathbf{0})\}_{k=1}^K$, with the mixing parameters being the multinomial probabilities determined by the gating network. We may thus write

$$\begin{aligned} f_D(d | \mathbf{x}, \mathbf{0}) &= \sum_{k=1}^K g_k f_D(d | \mathbf{x}, k, \mathbf{0}) \\ &= \frac{1}{\sqrt{2\pi}} \sum_{k=1}^K g_k \exp\left(-\frac{1}{2}(d - y_k)^2\right) \end{aligned} \quad (7.30)$$

The probability distribution of Eq. (7.30) is called an *associative Gaussian mixture model*. Its nonassociative counterpart is the traditional Gaussian mixture model (Titterton et al., 1985; McLachlan and Basford, 1988), which is briefly described in

Chapter 5. An associative model differs from a nonassociative model in that the conditional means μ_k and mixing parameters g_k are *not* fixed; rather, they are all functions of the input vector $\mathbf{x}$. The associative Gaussian mixture model of Eq. (7.30) may therefore be viewed as a generalization of the traditional Gaussian mixture model.

The important aspects of the ME model shown in Fig. 7.8, assuming that it is properly tuned through training, are:

1. The output y_k of the k th expert provides an estimate of the conditional mean of the random variable representing the desired response D , given $\mathbf{x}$ and that rule k of the probabilistic generative model holds.
2. The output g_k of the gating network defines the multinomial probability that the output of expert k matches the value $D = d$, on the basis of knowledge gained from $\mathbf{x}$ alone.

Working with the probability distribution of Eq. (7.30) and given the training sample $\{(\mathbf{x}_i, d_i)\}_{i=1}^N$, the problem is to *learn* the conditional means $\mu_k = y_k$ and the mixing parameters g_k , $k = 1, 2, \dots, K$, in an optimum manner, so that $f_D(d | \mathbf{x}, \mathbf{0})$ provides a good estimate of the underlying probability density function of the environment responsible for generating the training data.

Example 7.1 Regression Surface

Consider an ME model with two experts, and a gating network with two outputs denoted by g_1 and g_2 . The output g_1 is defined by (see Eq. (7.21))

$$\begin{aligned} g_1 &= \frac{\exp(u_1)}{\exp(u_1) + \exp(u_2)} \\ &= \frac{1}{1 + \exp(-(u_1 - u_2))} \end{aligned} \quad (7.31)$$

Let $\mathbf{a}_1$ and $\mathbf{a}_2$ denote the two weight vectors of the gating network. We may then write

$$u_k = \mathbf{x}^T \mathbf{a}_k, \quad k = 1, 2$$

and therefore rewrite Eq. (7.31) as:

$$g_1 = \frac{1}{1 + \exp(-\mathbf{x}^T(\mathbf{a}_1 - \mathbf{a}_2))} \quad (7.32)$$

The other output g_2 of the gating network is

$$\begin{aligned} g_2 &= 1 - g_1 \\ &= \frac{1}{1 + \exp(-\mathbf{x}^T(\mathbf{a}_2 - \mathbf{a}_1))} \end{aligned}$$

Thus, both g_1 and g_2 are in the form of a logistic function, but with a difference. The orientation of g_1 is determined by the direction of the difference vector $(\mathbf{a}_1 - \mathbf{a}_2)$, whereas the orientation of g_2 is determined by the direction of the difference vector $(\mathbf{a}_2 - \mathbf{a}_1)$, that is the negative of that for gate g_1 . Along the *ridge* defined by $\mathbf{a}_1 = \mathbf{a}_2$, we have $g_1 = g_2 = 1/2$, and the two experts contribute equally to the output of the ME model. Away from the ridge, one or the other of the two experts assumes the dominant role.

7.7 HIERARCHICAL MIXTURE OF EXPERTS MODEL

The ME model of Fig. 7.8 works by dividing the input space into different subspaces, with a single gating network responsible for distributing information (gathered from the training data) to the various experts. The *hierarchical mixtures of experts (HME) model*, illustrated in Fig. 7.11, is a natural extension of the ME model. The illustration is for an HME model of four experts. The architecture of the HME model is like a *tree*, in which the gating networks sit at the various nonterminals of the tree and the experts sit at the leaves of the tree. The HME model differs from the ME model in that the input space is divided into a *nested* set of subspaces, with the information combined and redistributed among the experts under the control of several gating networks arranged in a *hierarchical* manner.

The HME model of Fig. 7.11 has two *levels of hierarchy* or two *layers of gating networks*. By continuing with the application of the principle of divide and conquer in a manner similar to that illustrated, we may construct an HME model with any number of levels of hierarchy. Note that according to the convention described in Fig. 7.11, the numbering of gating levels starts from the output node of the tree.




FIGURE 7.11 Hierarchical mixture of experts (HME) illustrated for two levels of hierarchy.

The formulation of the HME model of Fig. 7.11 can be viewed in two ways (Jordan, 1994):

1. *The HME model is a product of the divide and conquer strategy.* If we believe that it is a good strategy to divide the input space into regions, then it is an equally good strategy to divide each of those regions into subregions. We can continue recursively in this manner until we reach a stage where the complexity of the approximating surfaces is a good fit to the "local" complexity of the training data. The HME model should therefore perform at least as well as, and often better than, the ME model for this reason: A higher level gating network in the HME model effectively combines information and redistributes it to the experts in the particular subtree controlled by that gating network. Consequently, each parameter in the subtree in question shares strength with other parameters contained in the same subtree, thereby contributing to possible improvement in the overall performance of the HME model.
2. *The HME model is a soft-decision tree.* According to this second viewpoint, the mixture of experts is just a one-level decision tree, sometimes referred to as a *decision stump*. In a more general setting, the HME model is viewed as the probabilistic framework for a decision tree, with the output node of the HME model referred to as the *root* of the tree. The methodology of a *standard* decision tree constructs a tree that leads to *hard* (e.g., yes-no) decisions in different regions of the input space. This is to be contrasted with the soft decisions performed by an HME model. Consequently, the HME model may outperform the standard decision tree for two reasons:
 - A hard decision inevitably results in loss of information, whereas a soft decision tries to preserve information. For example, a soft binary decision conveys information about distance from the decision boundary (i.e., the point at which the decision is 0.5), whereas a hard decision cannot. We may therefore say that unlike the standard decision tree, the HME model adheres to the *information preservation rule*. This empirical rule states that the information content of an input signal should be preserved in a computationally efficient manner until the system is ready for final decision-making or parameter estimation (Haykin, 1996).
 - Standard decision trees suffer from a *greediness* problem. Once a decision is made in such a tree, it is frozen and never changes thereafter. The HME model lessens the greediness problem because the decisions made throughout the tree are continually altered. Unlike the standard decision tree, it is possible in the HME model to recover from a poor decision somewhere further along the tree.

The second viewpoint, that is, a soft decision tree is the preferred way to think about an HME model. With the HME viewed as the probabilistic basis for a decision tree, it allows us to calculate a likelihood for any given data set, and maximize that likelihood with respect to the parameters that determine the splits between various regions in the input space. Thus, by building on what we already know about standard decision trees, we may have a practical solution to the model selection problem as discussed in the next section.

MODEL SELECTION USING A STANDARD DECISION TREE

As with every other neural network, a satisfactory solution to the parameter estimation problem hinges on the selection of a suitable model for the problem at hand. In the case of an HME model, the model selection involves the choice of the number and arrangement of the decision nodes in the tree. One practical solution to this particular model selection problem is to run a standard decision tree algorithm on the training data, and to adopt the tree so obtained as the initializing step for the learning algorithm used to determine the parameters of the HME model (Jordan, 1994).

The HME model has clear similarities with standard decision trees, such as the *classification and regression tree* (CART) due to Breiman et al., (1984). Figure 7.12 shows an example of CART, where the space of input data, $\mathcal{X}$ is repeatedly partitioned by a sequence of binary splits into *terminal nodes*. Comparing Fig. 7.12 with Fig. 7.11, we readily see the following similarities between CART and HME:

- The rules for selecting splits at intermediate (i.e., nonterminal) nodes of CART play an analogous role to gating networks in the HME model.
- The terminal nodes in CART play an analogous role to expert networks in the HME model.

By starting with CART for a classification or regression problem of interest, we take advantage of the *discrete* nature of CART to provide an efficient search among alternative trees. By using a tree so chosen as the initializing step in the learning algorithm for parameter estimation, we take advantage of the *continuous* probabilistic basis of the HME model to yield an improved “soft” estimate of the desired response.

The CART Algorithm


In light of what has just been said, a brief description of the CART algorithm is in order. The description is presented in the context of regression. Starting with the training data $\{(\mathbf{x}_i, d_i)\}_{i=1}^N$, we may use CART to construct a binary tree T for least-squares regression, by proceeding as follows (Breiman et al., 1984):

1. *Selection of splits.* Let a node t denote a subset of the current tree T . Let $\bar{d}(t)$ denote the average of d_i for all cases $(\mathbf{x}_i, d_i)$ falling into t , that is,

$$\bar{d}(t) = \frac{1}{N(t)} \sum_{\mathbf{x}_i \in t} d_i \quad (7.33)$$

FIGURE 7.12 Binary decision tree, described as follows:

- Nodes t_2 and t_3 are descendants of node t_1 .
- Nodes t_4 and t_5 are descendants of node t_2 ; and likewise for t_6 and t_7 in relation to t_3 .



where the sum is over all d_i such that $\mathbf{x}_i \in t$ and $N(t)$ is the total number of cases in t . Define

$$\mathcal{E}(t) = \frac{1}{N} \sum_{\mathbf{x}_i \in t} (d_i - \bar{d}(t))^2 \quad (7.34)$$

and

$$\mathcal{E}(T) = \sum_{t \in T} \mathcal{E}(t) \quad (7.35)$$

For node t , the sum $\sum_{\mathbf{x}_i \in t} (d_i - \bar{d}(t))^2$ represents the “within node sum of squares”; that is, it is the total squared deviations of all the d_i in t from their average $\bar{d}(t)$. Summing these deviations over $t \in T$ gives the total node sum of squares and dividing it by N gives the average.

Given any set of splits S of a current node t in T , the best split s^* is that split in S that most decreases $\mathcal{E}(T)$. To be more precise, suppose that for any split s of node t into t_L (a new node to the left of t) and t_R (another new node to the right of t), we let

$$\Delta \mathcal{E}(s, t) = \mathcal{E}(T) - \mathcal{E}(t_L) - \mathcal{E}(t_R) \quad (7.36)$$

The best split s^* is then taken to be the particular split for which we have

$$\Delta \mathcal{E}(s^*, t) = \max_{s \in S} \Delta \mathcal{E}(t, s) \quad (7.37)$$

A regression tree so constructed is designed to maximize the decrease in $\mathcal{E}(T)$.

2. *Determination of a terminal node.* A node t is declared a terminal node if this condition is satisfied:

$$\max_{s \in S} \Delta \mathcal{E}(s, t) < \beta \quad (7.38)$$

where β is a prescribed threshold.

3. *Least-squares estimation of a terminal node's parameters.* Let t denote a terminal node in the final binary tree T , and let $\mathbf{X}(t)$ denote the matrix composed of $\mathbf{x}_i \in t$. Let $\mathbf{d}(t)$ denote the corresponding vector composed of all the d_i in t . Define

$$\mathbf{w}(t) = \mathbf{X}^+(t) \mathbf{d}(t) \quad (7.39)$$

where $\mathbf{X}^+(t)$ is the pseudoinverse of matrix $\mathbf{X}(t)$. The use of $\mathbf{w}(t)$ yields a least-squares estimate of $d(t)$ at the output of terminal node t . Using the weights calculated from Eq. (7.39), the split selection problem is solved by looking for the least sum of squared residuals (errors) with respect to the regression surfaces, rather than with respect to the means.

Using CART to Initialize the HME Model

Suppose that the CART algorithm has been applied to a set of training data, resulting in a binary decision tree for this problem. We may describe a split produced by CART as a multidimensional surface defined by

$$\mathbf{a}^T \mathbf{x} + b = 0$$

where $\mathbf{x}$ is the input vector, $\mathbf{a}$ denotes a parameter vector, and b denotes a bias.

Consider next the corresponding situation in an HME model. From Example 7.1 we note that the regression surface produced by a gating network in a binary tree may be expressed as:

$$g = \frac{1}{1 + \exp(-(\mathbf{a}^T \mathbf{x} + b))} \quad (7.40)$$

which defines a split, particularly when $g = 1/2$. Let the weight vector (difference) $\mathbf{a}$ for this particular gating network be written as

$$\mathbf{a} = \|\mathbf{a}\| \cdot \frac{\mathbf{a}}{\|\mathbf{a}\|} \quad (7.41)$$

where $\|\mathbf{a}\|$ denotes the length (i.e., Euclidean norm) of $\mathbf{a}$, and $\mathbf{a}/\|\mathbf{a}\|$ is a normalized unit-length vector. Using Eq. (7.41) in (7.40), we may thus rewrite a parameterized split at a gating network as:

$$g = \frac{1}{1 + \exp\left(-\|\mathbf{a}\| \left(\left(\frac{\mathbf{a}}{\|\mathbf{a}\|} \right)^T \mathbf{x} + \frac{b}{\|\mathbf{a}\|} \right)\right)} \quad (7.42)$$

where we see that $\mathbf{a}/\|\mathbf{a}\|$ determines the *direction* of the split and $\|\mathbf{a}\|$ determines the sharpness of the split. From the discussion presented in Chapter 2, we observe that the length of vector $\mathbf{a}$ acts effectively as the *reciprocal of temperature*. The important point to note from Eq. (7.42) is that a gating network made up of a linear filter followed by a softmax form of nonlinearity is able to mimic a split in the style of CART. Moreover, we have an additional degree of freedom, namely, the length of parameter vector $\mathbf{a}$. In a standard decision tree, this additional parameter is irrelevant because a threshold (i.e., hard decision) is used to create a split. In contrast, the length of $\mathbf{a}$ has a profound influence on the split sharpness produced by a gating network in the HME model. Specifically, for a synaptic weight vector $\mathbf{a}$ of fixed direction, we may state that:

- when $\mathbf{a}$ is long (i.e., the temperature is low), the split is sharp, and
- when $\mathbf{a}$ is short (i.e., the temperature is high), the split is soft.

If, in the limit we have $\|\mathbf{a}\| = 0$, the split vanishes and $g = 1/2$ on both sides of the vanished (fictitious) split. The effect of setting $\|\mathbf{a}\| = 0$ is equivalent to pruning the nonterminal node from the tree, because the gating network in question is no longer splitting. In the very extreme case when $\|\mathbf{a}\|$ is small (i.e., the temperature is high) at every nonterminal node, the entire HME model acts like a single node; that is, the HME is reduced to a linear regression model (assuming linear experts). As the synaptic weight vectors of the gating networks start to grow in length, the HME starts to make (soft) splits, thereby enlarging the number of degrees of freedom available to the model.

We may thus initialize the HME by proceeding as follows:

1. Apply CART to the training data.
2. Set the synaptic weight vectors of the experts in the HME model equal to the least-squares estimates of the parameter vectors at the corresponding terminal nodes of the binary tree resulting from the application of CART.

3. For the gating networks:

- (a) set the synaptic weight vectors to point in directions that are orthogonal to the corresponding splits in the binary tree obtained from CART, and
- (b) set the lengths (i.e., Euclidean norms) of the synaptic weight vectors equal to *small* random vectors.

7.9 A PRIORI AND A POSTERIORI PROBABILITIES

The multinomial probabilities g_k and $g_{j|k}$ pertaining to the first-level and second-level gating networks, respectively, may be viewed as *a priori* probabilities, in the sense that their values are solely dependent on the input vector (stimulus) $\mathbf{x}$. In a corresponding way, we may define *a posteriori* probabilities $h_{j|k}$ and h_k whose values depend on both the input vector $\mathbf{x}$ and the responses of the experts to $\mathbf{x}$. This latter set of probabilities is useful in the development of learning algorithms for HME models.

Referring to the HME model of Fig. 7.11, we define the *a posteriori* probabilities at the nonterminal nodes of the tree as (Jordan and Jacobs, 1994):

$$h_k = \frac{g_k \sum_{j=1}^2 g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)}{\sum_{k=1}^2 g_k \sum_{j=1}^2 g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)} \quad (7.43)$$

and

$$h_{j|k} = \frac{g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)}{\sum_{j=1}^2 g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)} \quad (7.44)$$

The product of h_k and $h_{j|k}$ defines the *joint a posteriori* probability that expert (j, k) produces an output y_{jk} that matches the desired response d , as given by

$$h_{jk} = h_k h_{j|k} = \frac{g_k g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)}{\sum_{k=1}^2 g_k \sum_{j=1}^2 g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right)} \quad (7.45)$$

The probability h_{jk} satisfies the following two conditions

$$0 \leq h_{jk} \leq 1 \quad \text{for all } (j, k) \quad (7.46)$$

and

$$\sum_{j=1}^2 \sum_{k=1}^2 h_{jk} = 1 \quad (7.47)$$

The implication of Eq. (7.47) is that credit is distributed across the experts on a competitive basis. Moreover, we note from Eq. (7.45) that the closer y_{jk} is to d , the more

likely it is that expert (j, k) is given credit for its output having matched d , which is intuitively satisfying.

An important feature of the HME model that deserves special mention is the *recursiveness* in the computations involved in calculating the *a posteriori* probabilities. In examining Eqs. (7.43) and (7.44), we see that the denominator of $h_{j|k}$ in Eq. (7.44) appears in the numerator of h_k in Eq. (7.43). In an HME model we want to calculate the *a posteriori* probability for every nonterminal node in the tree. This is where recursiveness is of particular value. Specifically, the calculation of the *a posteriori* probabilities of all the nonterminal nodes in the tree is achieved in a single pass as described here:

- In moving through the tree toward the root node, level by level, the *a posteriori* probability at any nonterminal node of the tree is obtained simply by combining the *a posteriori* probabilities of its "children".

7.10 MAXIMUM LIKELIHOOD ESTIMATION

Turning next to the issue of parameter estimation for the HME model, we first note that its probabilistic interpretation is somewhat different from that of the ME model. With the HME model formulated as a binary tree, it is assumed that the environment responsible for generating the data involves a *nested sequence of soft (binary) decisions, terminating on the regression of the input vector $\mathbf{x}$ onto the output d* . In particular, we assume that in the *probabilistic generative model* for the HME, the decisions are modeled as multinomial random variables (Jordan and Jacobs, 1994). That is, for each input $\mathbf{x}$ we interpret $g_k(\mathbf{x}, \theta_k^0)$ as the multinomial probabilities associated with the first decision, and $g_{j|k}(\mathbf{x}, \theta_{jk}^0)$ as the conditional multinomial distributions associated with the second decision. As before, the superscript 0 signifies true values of the generative model parameters. The decisions form a decision tree. As with the ME model, softmax is used as the activation function of the gating networks throughout the HME model. In particular, the activation g_k of the k th output neuron in the *top-level gating network* is defined by

$$g_k = \frac{\exp(u_k)}{\exp(u_1) + \exp(u_2)}, \quad k = 1, 2 \quad (7.48)$$

where u_k is the weighted sum of the inputs applied to that neuron. Similarly, the activation of the j th output neuron in the k th gating network in the second level of hierarchy is defined by

$$g_{j|k} = \frac{\exp(u_{jk})}{\exp(u_{1k}) + \exp(u_{2k})}, \quad (j, k) = 1, 2 \quad (7.49)$$

where u_{jk} is the weighted sum of the inputs applied to this particular neuron.

For the sake of presentation we will work with an HME model with only two levels of hierarchy (i.e., two layers of gating networks), as indicated in Fig. 7.11. As with the ME model, each of the experts in the HME model is assumed to consist of a single

layer of linear neurons. Let y_{jk} denote the output of expert (j, k) . We may then express the overall output of the HME model as

$$y = \sum_{k=1}^2 g_k \sum_{j=1}^2 g_{j|k} y_{jk} \quad (7.50)$$

Following a procedure similar to that described for the ME model in Section 7.6 we may formulate the probability density function of the random variable D representing the desired response for the HME model of Fig. 7.11, given the input $\mathbf{x}$, as follows:

$$f_D(d|\mathbf{x}, \mathbf{0}) = \frac{1}{\sqrt{2\pi}} \sum_{k=1}^2 g_k \sum_{j=1}^2 g_{j|k} \exp\left(-\frac{1}{2}(d - y_{jk})^2\right) \quad (7.51)$$

Thus, for a given set of training data, Eq. (7.51) defines a model for the underlying distribution of the data. The vector $\mathbf{0}$ encompasses all the synaptic weights involved in the characterization of the gating and expert networks of the HME model.

The designation *likelihood function*, denoted by $l(\mathbf{0})$, is given to the probability density function $f_D(d|\mathbf{x}, \mathbf{0})$, viewed as a function of the parameter vector $\mathbf{0}$. We thus write

$$l(\mathbf{0}) = f_D(d|\mathbf{x}, \mathbf{0}) \quad (7.52)$$

Although the conditional joint probability density function and the likelihood function have exactly the same formula, it is vital that we appreciate the physical difference between them. In the case of $f_D(d|\mathbf{x}, \mathbf{0})$, the input vector $\mathbf{x}$ and parameter vector $\mathbf{0}$ are fixed but the desired response d is variable. However, in the case of the likelihood function $l(\mathbf{0})$, both $\mathbf{x}$ and d are fixed but $\mathbf{0}$ is variable.

In practice, we find it more convenient to work with the natural logarithm of the likelihood function rather than the likelihood itself. Using $L(\mathbf{0})$ to denote the *log-likelihood function*, we write

$$L(\mathbf{0}) = \log[l(\mathbf{0})] \\ = \log[f_D(d|\mathbf{x}, \mathbf{0})] \quad (7.53)$$

The logarithm of $l(\mathbf{0})$ is a monotonic transformation of $l(\mathbf{0})$. This means that whenever $l(\mathbf{0})$ increases, its logarithm $L(\mathbf{0})$ also increases. Since $l(\mathbf{0})$ is a formula for a conditional probability density function, it can never become negative. It follows therefore that there is no problem in evaluating $L(\mathbf{0})$. Hence, an estimate $\hat{\mathbf{0}}$ of parameter vector $\mathbf{0}$ can be obtained as a solution of the *likelihood equation*

$$\frac{\partial}{\partial \mathbf{0}} l(\mathbf{0}) = \mathbf{0}$$

or equivalently the *log-likelihood equation*

$$\frac{\partial}{\partial \mathbf{0}} L(\mathbf{0}) = \mathbf{0} \quad (7.54)$$

The term "maximum likelihood estimate" with the desired asymptotic properties⁵ usually refers to a root of the likelihood equation that globally maximizes

the likelihood function $L(\mathbf{0})$. The estimate $\hat{\mathbf{0}}$ used in practice, however, may in actual fact be a local maximum and not a global maximum. In any event, maximum likelihood estimation, due to Fisher (1925), is based on a relatively simple idea:

Different populations generate different data samples and any given data sample is more likely to have come from some population than from some others.

More specifically, the unknown parameter vector $\mathbf{0}$ is estimated by its *most plausible value*, given the input vector $\mathbf{x}$. In other words, the maximum likelihood estimate $\hat{\mathbf{0}}$ is that value of the parameter vector $\mathbf{0}$ for which the conditional probability density function $f_{\mathbf{0}}(d|\mathbf{x}, \mathbf{0})$ is maximum.

LEARNING STRATEGIES FOR THE HME MODEL

The probabilistic description of the HME model in Section 7.10 has led us to the log-likelihood function $L(\mathbf{0})$ as the objective function to be maximized. The key question is how to perform this maximization. As with every optimization problem, there is no unique approach to the maximization of $L(\mathbf{0})$. Rather, we have several approaches at our disposal, two of which are summarized here (Jacobs and Jordan, 1991; Jordan and Jacobs, 1994,):

1. *Stochastic gradient approach.* This approach yields an on-line algorithm for the maximization of $L(\mathbf{0})$. Its formulation for a two-level HME model, as depicted in Fig. 7.11, hinges on formulas for the following ingredients:

- The gradient vector $\partial L / \partial \mathbf{w}_{jk}$ for the vector of synaptic weights in expert (j, k) .
- The gradient vector $\partial L / \partial \mathbf{a}_k$ for the vector of synaptic weights in output neuron k of the top-level gating network.
- The gradient vector $\partial L / \partial \mathbf{a}_{jk}$ for the vector of synaptic weights in the output neuron of the second-level gating network associated with expert (j, k) .

It is a straightforward matter to show that (see Problem 7.9):

$$\frac{\partial L}{\partial \mathbf{w}_{jk}} = h_{j|k}(n) h_k(n)(d(n) - y_{jk}(n))\mathbf{x}(n) \quad (7.55)$$

$$\frac{\partial L}{\partial \mathbf{a}_k} = (h_k(n) - g_k(n))\mathbf{x}(n) \quad (7.56)$$

$$\frac{\partial L}{\partial \mathbf{a}_{jk}} = h_k(n)(h_{j|k}(n) - g_{j|k}(n))\mathbf{x}(n) \quad (7.57)$$

Equation (7.55) states that during the training process, the synaptic weights of expert (j, k) are adjusted to correct for the error between the output y_{jk} and the desired response d , in proportion to the joint *a posteriori* probability h_{jk} that expert (j, k) produces an output that matches d . Equation (7.56) states that the synaptic weights of output neuron k in the top-level gating network are adjusted so as to force the *a priori* probabilities $g_k(n)$ to move toward the corresponding *a posteriori* probabilities $h_k(n)$. Equation (7.57) states that the synaptic weights of the output neuron of the second-level gating network associated with expert (j, k) are adjusted to correct for the error

between the *a priori* probability $g_{j|k}$ and the corresponding *a posteriori* probability $h_{j|k}$ in proportion to the *a posteriori* probability $h_k(n)$.

According to Eqs. (7.55) to (7.57), the synaptic weights of the HME model are updated after the presentation of each pattern (stimulus). By summing the gradient vectors shown here over n , we may formulate the batch version of the gradient ascent algorithm for maximizing the log-likelihood function $L(\mathbf{0})$.

2. *Expectation-maximization approach.* The *expectation-maximization (EM) algorithm*, due to Dempster et al. (1977), provides an iterative procedure for computing maximum likelihood estimates in situations where, except for some missing data, the issue of maximum likelihood estimation would be a straightforward matter. The EM algorithm derives its name from the fact that on each iteration of the algorithm there are two steps:

- *Expectation step or E-step*, which uses the observed data set of an *incomplete data problem* and the current value of the parameter vector to manufacture data so as to postulate an augmented or so-called *complete data set*.
- *Maximization step or M-step*, which consists of deriving a new estimate of the parameter vector by maximizing the log-likelihood function of the complete data manufactured in the E-step.

Thus, starting from a suitable value for the parameter vector, the E-step and M-step are repeated on an alternate basis until convergence.

The situations where the EM algorithm is applicable include not only those that involve naturally incomplete data, but also a variety of other situations where the incompleteness of data is not at all evident in or natural to the problem of interest. Indeed, computation of the maximum likelihood estimate is often greatly facilitated by artificially formulating it as an incomplete data problem. This is done because the EM algorithm is able to exploit the reduced complexity of the maximum likelihood estimation, given the complete data (McLachlan and Krishnan, 1997). The HME model is one such example application. In this case, missing data in the form of certain indicator variables are artificially introduced into the HME model to facilitate the maximum likelihood estimation of the unknown parameter vector, as described in Section 7.12.

An important feature of the HME model, whether designed using the stochastic gradient approach or the EM algorithm, is two-fold:

- Each gating network in the model is continually computing the *a posteriori* probability for every data point in the training set.
- The adjustments applied to the synaptic weights of the expert and gating networks in the model, from one iteration to the next, are functions of the *a posteriori* probability thus computed and the corresponding *a priori* probability.

Accordingly, if an expert network lower down in the tree fails to do a good job in fitting the training data in its local neighborhood, the regression (discriminant) surface of a gating network higher up in the tree will be moved around. This movement can, in turn, help the experts on the next iteration of the learning algorithm to fit the data better by shifting the subspaces in which they are supposed to do their data fitting. This is the process by which the HME model is able to ameliorate the greediness problem inherent to a standard decision tree like CART.

12 EM ALGORITHM

The EM algorithm is remarkable in part because of the simplicity and generality of the underlying theory, and in part because of the wide range of applications that fall under its umbrella.⁶ In this section we present a description of the EM algorithm in a generic sense. We go on to consider its application to the parameter estimation problem in the HME model in the next section.

Let the vector $\mathbf{z}$ denote the missing or unobservable data. Let $\mathbf{r}$ denote the complete-data vector, made up of some observable data d and the missing data vector $\mathbf{z}$. There are therefore two data spaces, $\mathcal{R}$ and $\mathcal{D}$, to be considered, with the mapping from $\mathcal{R}$ to $\mathcal{D}$ being many-to-one. However, instead of observing the complete data vector $\mathbf{r}$, we are actually only able to observe the incomplete data $d = d(\mathbf{r})$ in $\mathcal{D}$.

Let $f_c(\mathbf{r}|\mathbf{0})$ denote the conditional pdf of $\mathbf{r}$, given a parameter vector $\mathbf{0}$. It follows therefore that the conditional pdf of random variable D , given $\mathbf{0}$, is defined by

$$f_D(d|\mathbf{0}) = \int_{\mathcal{R}(d)} f_c(\mathbf{r}|\mathbf{0}) d\mathbf{r} \quad (7.58)$$

where $\mathcal{R}(d)$ is the subspace of $\mathcal{R}$ that is determined by $d = d(\mathbf{r})$. The EM algorithm is directed at finding a value of $\mathbf{0}$ that maximizes the *incomplete-data log-likelihood function*

$$L(\mathbf{0}) = \log f_D(d|\mathbf{0})$$

This problem, however, is solved indirectly by working *iteratively* with the *complete-data log-likelihood function*

$$L_c(\mathbf{0}) = \log f_c(\mathbf{r}|\mathbf{0}) \quad (7.59)$$

which is a random variable, because the missing data vector $\mathbf{z}$ is unknown.

To be more specific, let $\hat{\mathbf{0}}(n)$ denote the value of the parameter vector $\mathbf{0}$ on iteration n of the EM algorithm. In the E-step of this iteration, we calculate the expectation

$$Q(\mathbf{0}, \hat{\mathbf{0}}(n)) = E[L_c(\mathbf{0})] \quad (7.60)$$

where the expectation is performed with respect to $\hat{\mathbf{0}}(n)$. In the M-step of this same iteration, we maximize $Q(\mathbf{0}, \hat{\mathbf{0}}(n))$ with respect to $\mathbf{0}$ over the parameter (weight) space $\mathcal{W}$, and so find the updated parameter estimate $\hat{\mathbf{0}}(n+1)$, as shown by

$$\hat{\mathbf{0}}(n+1) = \arg \max_{\mathbf{0}} Q(\mathbf{0}, \hat{\mathbf{0}}(n)) \quad (7.61)$$

The algorithm is started with some initial value $\hat{\mathbf{0}}(0)$ of the parameter vector $\mathbf{0}$. The E-step and M-step are then alternately repeated in accordance with Eqs. (7.60) and (7.61), respectively, until the difference between $L(\hat{\mathbf{0}}(n+1))$ and $L(\hat{\mathbf{0}}(n))$ drops to some arbitrary small value; at that point the computation is terminated.

Note that after an iteration of the EM algorithm, the incomplete-data log-likelihood function is *not* decreased, as shown by (see Problem 7.10)

$$L(\hat{\mathbf{0}}(n+1)) \geq L(\hat{\mathbf{0}}(n)) \quad \text{for } n = 0, 1, 2, \dots, \quad (7.62)$$

Equality usually means that we are at a stationary point of the log-likelihood function.⁷

7.13 APPLICATION OF THE EM ALGORITHM TO THE HME MODEL

Having familiarized ourselves with the EM algorithm, we are now ready to solve the parameter estimation problem in the HME model using the EM algorithm.⁸

Let $g_k^{(i)}$ and $g_{jk}^{(i)}$ denote the (conditional) multinomial probabilities associated with the decisions taken by the first-level gating network k and second-level gating network (j, k) of the HME model in Fig. 7.11, respectively, when it operates under example i of the training set. Then, from Eq. (7.51) we readily see that the corresponding value of the conditional pdf of the random variable D , given example $\mathbf{x}_i$ and parameter vector $\mathbf{0}$, is given by

$$f_D(d_i|\mathbf{x}_i, \mathbf{0}) = \frac{1}{\sqrt{2\pi}} \sum_{k=1}^2 g_k^{(i)} \sum_{j=1}^2 g_{jk}^{(i)} \exp\left(-\frac{1}{2}(d_i - y_{jk}^{(i)})^2\right) \quad (7.63)$$

where $y_{jk}^{(i)}$ is the output produced by expert (j, k) in response to the i th example of the training set. Assuming that all the N examples contained in the training set are statistically independent, we may formulate the log-likelihood function for the incomplete data problem as follows:

$$L(\mathbf{0}) = \log \left[\prod_{i=1}^N f_D(d_i|\mathbf{x}_i, \mathbf{0}) \right] \quad (7.64)$$

Using Eq. (7.63) in (7.64) and ignoring the constant $-(1/2)\log(2\pi)$, we may write

$$L(\mathbf{0}) = \sum_{i=1}^N \log \left[\sum_{k=1}^2 g_k^{(i)} \sum_{j=1}^2 g_{jk}^{(i)} \exp\left(-\frac{1}{2}(d_i - y_{jk}^{(i)})^2\right) \right] \quad (7.65)$$

To compute the maximum likelihood estimate of $\mathbf{0}$, we have to find a stationary point (i.e., local or global maximum) of $L(\mathbf{0})$. Unfortunately, the log-likelihood function $L(\mathbf{0})$, as defined in Eq. (7.65), does not lend itself readily to this kind of computation.

To overcome this computational difficulty, we *artificially* augment the observable data $\{d_i\}_{i=1}^N$ by including a corresponding set of missing data in accordance with the EM algorithm. We do so by introducing *indicator variables* that pertain to the probability model of the HME architecture as follows (Jordan and Jacobs, 1994):

- $z_k^{(i)}$ and $z_{jk}^{(i)}$ are interpreted as the labels that correspond to decisions made in the probability model for the i th example in the training set. These variables are defined in such a way that only a single one of the $z_k^{(i)}$ is equal to one and only a single one of the $z_{jk}^{(i)}$ is equal to one for all i . Both $z_k^{(i)}$ and $z_{jk}^{(i)}$ are treated as statistically independent discrete random variables with their respective expectations defined by

$$E[z_k^{(i)}] = P[z_k^{(i)} = 1 | \mathbf{x}_i, d_i, \hat{\mathbf{0}}(n)] = h_k^{(i)} \quad (7.66)$$

and

$$E[z_{jk}^{(i)}] = P[z_{jk}^{(i)} = 1 | \mathbf{x}_i, d_i, \hat{\mathbf{0}}(n)] = h_{jk}^{(i)} \quad (7.67)$$

where $\hat{\mathbf{0}}(n)$ is the estimate of parameter vector $\mathbf{0}$ at iteration n of the EM algorithm.

- $z_{jk}^{(i)} = z_{j|k}^{(i)} z_k^{(i)}$ is interpreted as the label that specifies the expert (j, k) in the probability model for the i th example in the training sample. It is also treated as a discrete random variable with its expectation, defined by

$$\begin{aligned} E[z_{jk}^{(i)}] &= E[z_{j|k}^{(i)} z_k^{(i)}] \\ &= E[z_{j|k}^{(i)}] E[z_k^{(i)}] \\ &= h_{j|k}^{(i)} h_k^{(i)} = h_{jk}^{(i)} \end{aligned} \quad (7.68)$$

The $h_k^{(i)}$, $h_{j|k}^{(i)}$ and $h_{jk}^{(i)}$ in Eqs. (7.66) to (7.68) are the *a posteriori* probabilities introduced in Section 7.9; the superscript i has been added to them to signify the training example in question. See Problem 7.13 for a justification of these three equations.

By adding the missing data thus defined to the observable data, the maximum likelihood estimation problem is considerably simplified. More specifically, let $f_c(d_i, z_{jk}^{(i)} | \mathbf{x}_i, \mathbf{0})$ denote the conditional pdf of the complete data made up of d_i and $z_{jk}^{(i)}$, given $\mathbf{x}_i$ and parameter vector $\mathbf{0}$. We then write

$$f_c(d_i, z_{jk}^{(i)} | \mathbf{x}_i, \mathbf{0}) = \prod_{j=1}^2 \prod_{k=1}^2 (g_k^{(i)} g_{j|k}^{(i)} f_{jk}(d_i)) \quad (7.69)$$

where $f_{jk}(d_i)$ is the conditional pdf of d_i , given that expert (j, k) in the HME model is chosen; that is, $f_{jk}(i)$ is given by the Gaussian distribution:

$$f_{jk}(d_i) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} (d_i - y_{jk}^{(i)})^2\right) \quad (7.70)$$

Note that the formula of Eq. (7.69) corresponds to a hypothetical experiment, containing indicator variables represented by $z_{jk}^{(i)}$ that are nonobservable in a physical data sense. In any event, the log-likelihood function for the complete-data problem, accounting for the entire training set, is given by

$$\begin{aligned} L_c(\mathbf{0}) &= \log \left[\prod_{i=1}^N f_c(d_i, z_{jk}^{(i)} | \mathbf{x}_i, \mathbf{0}) \right] \\ &= \log \left[\prod_{i=1}^N \prod_{j=1}^2 \prod_{k=1}^2 (g_k^{(i)} g_{j|k}^{(i)} f_{jk}(d_i))^{z_{jk}^{(i)}} \right] \\ &= \sum_{i=1}^N \sum_{j=1}^2 \sum_{k=1}^2 z_{jk}^{(i)} [\log g_k^{(i)} + \log g_{j|k}^{(i)} + \log f_{jk}(d_i)] \end{aligned} \quad (7.71)$$

Using Eq. (7.70) in (7.71) and ignoring the constant $-(1/2)\log(2\pi)$, we may therefore write

$$L_c(\mathbf{0}) = \sum_{i=1}^N \sum_{j=1}^2 \sum_{k=1}^2 z_{jk}^{(i)} \left[\log g_k^{(i)} + \log g_{j|k}^{(i)} - \frac{1}{2} (d_i - y_{jk}^{(i)})^2 \right] \quad (7.72)$$

Comparing Eq. (7.72) with (7.65), we immediately see the computational benefit gained by adding the indicator variables as missing data to the set of observables: The

maximum likelihood estimation problem has been decoupled into a set of regression problems for the individual experts and a separate set of multinomial classification problems for the gating networks.

To proceed with the application of the EM algorithm, we first invoke the E-step of the algorithm by taking the expectation of the complete data log-likelihood function $L_c(\mathbf{0})$, as shown by

$$\begin{aligned} Q(\mathbf{0}, \hat{\mathbf{0}}(n)) &= E[L_c(\mathbf{0})] \\ &= \sum_{i=1}^N \sum_{j=1}^2 \sum_{k=1}^2 E[z_{jk}^{(i)}] \cdot \left(\log g_k^{(i)} + \log g_{j|k}^{(i)} - \frac{1}{2} (d_i - y_{jk}^{(i)})^2 \right) \end{aligned} \quad (7.73)$$

where the expectation operator is shown acting on the indicator variable $z_{jk}^{(i)}$ as this is the only unobservable variable. Hence, using Eq. (7.68) in (7.73), we obtain (Jordan and Jacobs, 1994):

$$Q(\mathbf{0}, \hat{\mathbf{0}}(n)) = \sum_{i=1}^N \sum_{j=1}^2 \sum_{k=1}^2 h_{jk}^{(i)} \left(\log g_k^{(i)} + \log g_{j|k}^{(i)} - \frac{1}{2} (d_i - y_{jk}^{(i)})^2 \right) \quad (7.74)$$

The M-step of the algorithm requires maximizing $Q(\mathbf{0}, \hat{\mathbf{0}}(n))$ with respect to $\mathbf{0}$. The parameter vector $\mathbf{0}$ is made up of two sets of synaptic weights: one belonging to the gating networks and the other belonging to the experts. From our earlier discussion we note the following:

- The synaptic weights of the experts determine $y_{jk}^{(i)}$, which also enters the definition of $h_{jk}^{(i)}$. The expression $Q(\mathbf{0}, \hat{\mathbf{0}}(n))$ is therefore influenced by the experts only through the term $h_{jk}^{(i)} (d_i - y_{jk}^{(i)})^2$.
- The synaptic weights of the gating networks determine the probabilities $g_{j|k}^{(i)}$, $g_k^{(i)}$ and $h_{jk}^{(i)}$. The expression $Q(\mathbf{0}, \hat{\mathbf{0}}(n))$ is therefore influenced by the gating networks only through the term $h_{jk}^{(i)} (\log g_k^{(i)} + \log g_{j|k}^{(i)})$.

Accordingly, the M-step of the algorithm reduces to the following three optimization problems for an HME with two levels of hierarchy:

$$\mathbf{w}_{jk}(n+1) = \arg \min_{\mathbf{w}_{jk}} \sum_{i=1}^N h_{jk}^{(i)} (d_i - y_{jk}^{(i)})^2 \quad (7.75)$$

$$\mathbf{a}_j(n+1) = \arg \max_{\mathbf{a}_j} \sum_{i=1}^N \sum_{k=1}^2 h_{jk}^{(i)} \log g_k^{(i)} \quad (7.76)$$

and

$$\mathbf{a}_{jk}(n+1) = \arg \max_{\mathbf{a}_{jk}} \sum_{i=1}^N \sum_{l=1}^2 h_{ml}^{(i)} \sum_{m=1}^2 h_{ml}^{(i)} \log g_{ml}^{(i)} \quad (7.77)$$

The optimizations in Eqs. (7.75) to (7.77) are done with h fixed; h is a function of the parameters but the derivatives are not taken through h . Note also that all the quantities in the right-hand sides of these equations refer to measurements made at step n .

The optimization in Eq. (7.75), pertaining to the experts, is a weighted least-squares estimation problem. The remaining two optimization in Eqs. (7.76) and (7.77), pertaining to the gating networks, are maximum likelihood estimation problems.⁹ Note

also that although the equations are formulated for two levels of hierarchy, they may be readily extended to an arbitrary number of levels of hierarchy.

4 SUMMARY AND DISCUSSION

In the study of modeling, pattern classification, and regression problems, we have two extreme cases to consider:

1. *Simple models*, which provide insight into the problem of interest but lack accuracy.
2. *Complex models*, which provide accurate results but lack insight.

It is perhaps impossible to combine simplicity and accuracy in a single model. In the context of the discussion presented in the second part of this chapter, CART is an example of a simple model that uses hard decisions to partition the input space into a piecewise set of subspaces, with each subspace having its own expert. Unfortunately, the use of hard decisions results in some information loss and therefore loss in performance. The multilayer perceptron (MLP), on the other hand, is a complex model with a nested form of nonlinearity designed to preserve the information content of the training data. However, it uses a black box approach to globally fit a single function into the data, thereby losing insight into the problem. The HME, representing a dynamic type of committee machine, is a compromise model between these two extremes, sharing common features with both CART and MLP:

- The architecture of HME is similar to that of CART, but differs from it in soft partitioning rather than hard partitioning the input space.
- The HME uses a nested form of nonlinearity similar to that of MLP, not for the purpose of input-output mapping, but rather for partitioning the input space.

In this chapter we emphasize the use of two tools for the design of an HME model:

- CART as the architectural basis for dealing with the model selection problem.
- EM algorithm for solving the parameter estimation problem by iteratively computing maximum likelihood estimates of the model parameters.

The EM algorithm is usually guaranteed to move *uphill* in likelihood. Therefore, by using CART to initialize the EM algorithm in the manner described in Section 7.8, we may expect the EM algorithm to yield a better generalization performance than would be possible with the initial condition established by CART.

The EM algorithm is important and fundamental if the application of interest is that of maximum likelihood estimation, such as in *modeling*. An interesting modeling application is described in Jacobs, Jordan, and Barto (1991b), where an ME model is trained to do the what/where task. In this task, the model is required to determine what an object is, and where it is in the visual field. Two experts were used in the study, each of which is specialized for one aspect of the task. For a specific input, both experts generate outputs. But, the gating network decides the appropriate mixture for that input. The successful results reported by Jacobs et al. demonstrate that it is possible for a task assignment to be innately determined, not on the basis of the task *per se*, but rather by the match between a task's requirements and the computational properties of the model (Elman et al., 1996).

We finish this discussion by returning to the other class of committee machines studied in the first part of the chapter. Whereas the ME model or HME model relies on the use of gating network(s) activated by the input signal for fusing the knowledge acquired by the various experts in the model, a committee machine, based on the use of ensemble averaging or, alternatively, boosting, relies on the learning algorithm itself to do the integration, as summarized here:

1. Ensemble averaging improves error performance in a clever way by the combined use of two effects:
 - Reduction of error due to bias by purposely overfitting the individual experts in the committee machine.
 - Reduction of error due to variance by using different initial conditions in the training of the individual experts, and then ensemble-averaging their respective outputs.
2. Boosting improves error performance in an ingenious way of its own. In this case, individual experts are required to perform only slightly better than random guessing. The weak learnability of the experts is converted into strong learnability whereby the committee machine's error is made arbitrarily small. This remarkable conversion is achieved by *filtering* the distribution of the input data in a manner causing the weak learning models (i.e., experts) to eventually learn the entire distribution, or by *resampling* the training examples according to a certain probability distribution as in AdaBoost. The advantage of AdaBoost over boosting by filtering is that it works with a training sample of fixed size.

NOTES AND REFERENCES

1. *Ensemble-averaging methods* are discussed in Perrone (1993), where a large bibliography on the subject is included. Other references on this subject include Wolpert (1992), and Hashem (1997).
2. The use of ensemble-averaging for the design of a committee machine over a set of different initial conditions has been suggested by several neural network practitioners. However, the statistical analysis presented in Naftaly et al. (1997) and the procedure described therein for training a committee machine designed by ensemble-averaging over the space of initial conditions appear to be the first of their kind. In that paper, experimental results are presented based on the sunspot data and energy-prediction competition data. In both cases, significant reduction in variance is demonstrated by averaging over the space of initial conditions.
According to Naftaly et al. (1997), the use of popular training constraints such as weight decay and early stopping is *not* recommended in the design of a committee machine by ensemble averaging over the space of initial conditions.
3. The main references on boosting theory and related experimental studies, more or less in chronological order, are as follows: Schapire (1990), Drucker et al. (1993, 1994), Freund (1995), Breiman (1996b), Freund and Schapire (1996a, 1996b, 1997), Schapire (1997), and Schapire et al. (1997). The first references on the three basic approaches to boosting are as follows:
 - Filtering: Schapire (1990)
 - Resampling: Freund and Schapire (1996a)
 - Reweighting: Freund (1995)

4. The idea of using a mixture of experts for realizing a complex mapping function was first discussed by Jacobs, Jordan, Nowlan, and Hinton in their 1991a paper. The development of this model was motivated by (1) a proposal described in Nowlan (1990), viewing competitive adaptation in unsupervised learning as an attempt to fit a mixture of simple probability distributions (such as Gaussians) into a set of data points, and (2) ideas developed in the Ph.D. thesis of Jacobs (1990) using a similar modular architecture but a different cost function.
5. Maximum likelihood estimators have some desirable properties. Under quite general conditions, the following *asymptotic* properties may be proved (Kmenta, 1971):
 - (i) *Maximum likelihood estimators are consistent.* Let $L(\mathbf{0})$ denote the log-likelihood function and $\hat{\mathbf{0}}_n$ denote an element of the parameter vector $\mathbf{0}$. The partial derivative $\partial L / \partial \theta_i$ is called a *score*. We say that a maximum likelihood estimator is consistent in the sense that the value of $\hat{\mathbf{0}}_n$ for which the score $\partial L / \partial \theta_i$ is identically zero, *converges in probability* to the true value of θ_i as the sample size used in the estimation approaches infinity.
 - (ii) *Maximum likelihood estimators are asymptotically efficient.* That is,

$$\lim_{N \rightarrow \infty} \left\{ \frac{\text{var}[\hat{\mathbf{0}}_n - \hat{\mathbf{0}}_n]}{I_{ii}} \right\} = 1 \quad \text{for all } i$$

where N is the sample size, $\hat{\mathbf{0}}_n$ is the maximum likelihood estimate of $\mathbf{0}$, and I_{ii} is the i th diagonal element of the inverse of *Fisher's information matrix*. Fisher's information matrix is defined by

$$\mathbf{J} = - \begin{bmatrix} E \left[\frac{\partial^2 L}{\partial \theta_1^2} \right] & E \left[\frac{\partial^2 L}{\partial \theta_1 \partial \theta_2} \right] & \cdots & E \left[\frac{\partial^2 L}{\partial \theta_1 \partial \theta_M} \right] \\ E \left[\frac{\partial^2 L}{\partial \theta_2 \partial \theta_1} \right] & E \left[\frac{\partial^2 L}{\partial \theta_2^2} \right] & \cdots & E \left[\frac{\partial^2 L}{\partial \theta_2 \partial \theta_M} \right] \\ \vdots & \vdots & \ddots & \vdots \\ E \left[\frac{\partial^2 L}{\partial \theta_M \partial \theta_1} \right] & E \left[\frac{\partial^2 L}{\partial \theta_M \partial \theta_2} \right] & \cdots & E \left[\frac{\partial^2 L}{\partial \theta_M^2} \right] \end{bmatrix}$$

where M is the dimension of parameter vector $\mathbf{0}$.

- (iii) *Maximum likelihood estimators are asymptotically Gaussian.* That is, as the sample size approaches infinity, each element of the maximum likelihood estimate $\hat{\mathbf{0}}$ assumes a Gaussian distribution.
- In practice, we find the large sample (i.e., asymptotic) properties of maximum likelihood estimators hold rather well for sample size $N \geq 50$.
6. The paper by Newcomb (1886), considering the estimation of parameters of a mixture of two univariate Gaussian distributions, appears to be the earliest reference to an EM-type of process reported in the literature.

The name "EM algorithm" was coined by Dempster, Laird, and Rubin in their 1977 fundamental paper. In that paper, formulation of the EM algorithm for computing maximum likelihood estimates from incomplete data at various levels of generality was presented for the first time.

The first unified account of the theory, methodology, and applications of the EM algorithm, its history, and extensions was presented in book form by McLachlan and Krishnan (1997).

7. Under fairly general conditions, the likelihood values computed by the EM algorithm converge to stationary values. Wu (1983) presents a detailed account of the convergence

properties of the EM algorithm. However, the EM algorithm will *not* always lead to a local or a global maximum of the likelihood function. In Chapter 3 of the book by McLachlan and Krishnan (1997), two examples are presented where this is not the case. In one example the algorithm converges to a saddle point, and in the other example the algorithm converges to a local *minimum* of the likelihood function.

8. The EM algorithm can also handle Bayesian *maximum a posteriori (MAP) estimation* by incorporating *prior* information on the parameter vector; see Problem 7.11. Using Bayes' rule, we may express the conditional probability density function for parameter vector $\mathbf{0}$, given a set of observations $\mathbf{x}$, as

$$f_{\mathbf{0}}(\mathbf{0}|\mathbf{x}) = \frac{f_{\mathbf{x}}(\mathbf{x}|\mathbf{0})f_{\mathbf{0}}(\mathbf{0})}{f_{\mathbf{x}}(\mathbf{x})}$$

From this relation, we readily see that maximizing the *a posteriori* density $f_{\mathbf{0}}(\mathbf{0}|\mathbf{x})$ is equivalent to maximizing the product $f_{\mathbf{x}}(\mathbf{x}|\mathbf{0})f_{\mathbf{0}}(\mathbf{0})$, since $f_{\mathbf{x}}(\mathbf{x})$ is independent of $\mathbf{0}$. The probability density function $f_{\mathbf{0}}(\mathbf{0})$ represents the prior information available on $\mathbf{0}$. Maximizing $f_{\mathbf{0}}(\mathbf{0}|\mathbf{x})$ provides the *most probable* estimate of the parameter vector $\mathbf{0}$, given $\mathbf{x}$. Two points are noteworthy in the context of this estimate:

- Maximum likelihood estimation, represented by maximizing $f_{\mathbf{x}}(\mathbf{x}|\mathbf{0})$ with respect to $\mathbf{0}$, is a reduced form of maximum *a posteriori* estimation, reduced in the sense that it is void of prior information.
- The use of prior information is synonymous with regularization, which (we recall from Chapter 5) corresponds to a smooth input-output mapping.

In Waterhouse et al. (1996), a Bayesian framework for estimating parameters of a mixture of experts model is presented. The Bayesian approach described therein overcomes a phenomenon known as "overfitting", which leads to an estimate with high variance when using maximum likelihood inference.

9. An efficient algorithm, known as the *iteratively reweighted least-squares (IRLS) algorithm*, is available to solve the maximum likelihood estimation problems described in Eqs. (7.76) and (7.77); for a description of the IRLS algorithm, see McCullagh and Nelder (1989).

PROBLEMS

Ensemble Averaging

- 7.1 Consider a committee machine consisting of K experts. The input-output function of the k th expert is denoted by $F_k(\mathbf{x})$, where $\mathbf{x}$ is the input vector and $k = 1, 2, \dots, K$. The individual outputs of the experts are linearly combined to form the overall output y , defined by

$$y = \sum_{k=1}^K w_k F_k(\mathbf{x})$$

where w_k is a linear weight assigned to $F_k(\mathbf{x})$. The requirement is to evaluate w_k so that y provides a least-squares estimate of the desired response d corresponding to $\mathbf{x}$. Given a set of training data $\{(\mathbf{x}_i, d_i)\}_{i=1}^N$, determine the required values of the w_k 's to solve this parameter estimation problem.

Boosting

- 7.2 Compare the computational advantages and disadvantages of boosting by filtering and AdaBoost.
- 7.3 Ordinarily, boosting performs best on weak learning models, that is, learning models with relatively low generalization error rates. Suppose, however, you are given a strong learning

model, that is, a learning model with high generalization error rate. Assume that you are dealing with a training sample of fixed size. How do boosting by filtering and AdaBoost cope with this situation?

Mixture of Experts

7.4 Consider a piecewise-linear task described by

$$F(x_1, x_2, \dots, x_{10}) = \begin{cases} 3x_2 + 2x_3 + x_4 + 3 + \epsilon & \text{if } x_1 = 1 \\ 3x_5 + 2x_6 + x_7 - 3 + \epsilon & \text{if } x_1 = -1 \end{cases}$$

For comparison, the following network configurations are used:

1. Multilayer perceptron: "10→10→1" network
2. Mixture of experts: Gating networks: 10→2;
Expert networks: 10→1

Compare the computational complexities of these two networks.

- 7.5 The ME model described by the conditional probability density function of Eq. (7.30) is based on a scalar regression model, in which the error is Gaussian distributed with zero mean and unit variance.
- (a) Reformulate this equation for the more general case of an ME model corresponding to a multiple regression model, in which the desired response is a vector with dimension q and the error is a multivariate Gaussian distribution with zero mean and covariance matrix Σ .
 - (b) How is the ME model for this reformulation different from the ME model shown in Fig. 7.8?

7.6 Derive the stochastic gradient algorithm for the training of the mixture of experts model.

Hierarchical Mixture of Experts

- 7.7 (a) Construct the block diagram of an HME model with three levels of hierarchy. Assume the use of a binary decision tree for the model.
- (b) Write the *a posteriori* probabilities for the nonterminal nodes of the HME described in part (a). Demonstrate the recursiveness of the computations involved in evaluating these probabilities.
 - (c) Formulate the conditional probability density function for the HME model described in part (a).
- 7.8 Discuss the similarities and differences between HME models and radial-basis function (RBF) networks.
- 7.9 Derive the equations that describe the stochastic gradient algorithm for the training of an HME model with two levels of hierarchy. Assume a binary decision tree for the model.

EM Algorithm and its Application to the HME Model

- 7.10 Prove the monotonic-increasing property of the EM algorithm described in Eq. (7.62). For this derivation, do the following:
- (a) Let

$$k(\mathbf{r}|d, \boldsymbol{\theta}) = \frac{f_c(\mathbf{r}|\boldsymbol{\theta})}{f_D(d|\boldsymbol{\theta})}$$

denote the conditional probability density function of the augmented data vector $\mathbf{r}$, given the observation d and parameter vector $\boldsymbol{\theta}$. Hence, the incomplete data log-likelihood function may be expressed as

$$L(\boldsymbol{\theta}) = L_c(\boldsymbol{\theta}) - \log k(\mathbf{r}|d, \boldsymbol{\theta})$$

where $L_c(\boldsymbol{\theta}) = \log f_c(\mathbf{r}|\boldsymbol{\theta})$ is the complete data log-likelihood function. By taking the expectation of $L(\boldsymbol{\theta})$ with respect to the conditional distribution of $\mathbf{r}$, given d , show that

$$L(\boldsymbol{\theta}) = Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(n)) - K(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(n))$$

where

$$K(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(n)) = E[\log k(\mathbf{r}|d, \hat{\boldsymbol{\theta}})]$$

Hence, show that

$$L(\hat{\boldsymbol{\theta}}(n+1)) - L(\hat{\boldsymbol{\theta}}(n)) = [Q(\hat{\boldsymbol{\theta}}(n+1), \hat{\boldsymbol{\theta}}(n)) - Q(\hat{\boldsymbol{\theta}}(n), \hat{\boldsymbol{\theta}}(n))] \\ - [K(\hat{\boldsymbol{\theta}}(n+1), \hat{\boldsymbol{\theta}}(n)) - K(\hat{\boldsymbol{\theta}}(n), \hat{\boldsymbol{\theta}}(n))]$$

- (b) Jensen's inequality states that if $f(\cdot)$ is a convex function and u is a random variable, then

$$E[g(u)] \geq g(E[u])$$

where E is the expectation operator; moreover, if $g(\cdot)$ is strictly convex, then equality in this relation implies that $u = E[u]$ with probability 1 (Cover and Thomas, 1991). Using Jensen's inequality, show that

$$K(\hat{\boldsymbol{\theta}}(n+1), \hat{\boldsymbol{\theta}}(n)) - K(\hat{\boldsymbol{\theta}}(n), \hat{\boldsymbol{\theta}}(n)) \leq 0$$

Hence, show that Eq. (7.62) holds for $n = 0, 1, 2, \dots$

- 7.11 The EM algorithm is easily modified to accommodate the maximum *a posteriori* (MAP) estimate of a parameter vector $\boldsymbol{\theta}$. Using Bayes' rule, modify the E-step and M-step of the EM algorithm to provide for this estimation.
- 7.12 For an HME trained with the EM algorithm and an MLP trained with the back-propagation algorithm to provide a similar level of performance for a given task, we would intuitively expect the computational complexity of the HME to exceed that of the MLP. Argue in favor of or against the plausibility of this statement.
- 7.13 Justify the relations between the indicator variables and corresponding *a posteriori* probabilities described in Eqs. (7.66) to (7.68).
- 7.14 Equation (7.75) describes the weighted least-squares for the optimization of the expert networks in the HME model of Fig. 7.11, assuming that the desired response d is scalar. How is this relation modified for the case of a multidimensional desired response?

Principal Components Analysis

1 INTRODUCTION

An important feature of neural networks is the ability they have to *learn* from their environment, and through learning to *improve* performance in some sense. In the previous four chapters, the focus was on algorithms for supervised learning, for which a set of targets of interest is provided by an external teacher. The targets take the form of a desired input–output mapping, which the network is required to approximate. In this chapter and the next three chapters, we study algorithms for *self-organized learning* or *unsupervised learning*. The purpose of an algorithm for self-organized learning is to *discover* significant patterns or features in the input data, and to do the discovery *without* a teacher. To do so, the algorithm is provided with a set of rules of a *local* nature, which enables it to learn to compute an input–output mapping with specific desirable properties; the term “local” means that the change applied to the synaptic weight of a neuron is confined to the immediate neighborhood of that neuron. The modeling of network structures used for self-organized learning tends to follow neurobiological structures to a much greater extent than for supervised learning. This may not be surprising, because the process of network organization is fundamental to the organization of the brain.

The structure of a self-organizing system may take on a variety of different forms. It may, for example, consist of an *input (source) layer* and an *output (representation) layer*, with feedforward connections from input to output and lateral connections between neurons in the output layer. Another example is a feedforward network with multiple layers, in which the self-organization proceeds on a layer-by-layer basis. In both examples, the learning process consists of repeatedly modifying the synaptic weights of all the connections in the system in response to input (activation) patterns and in accordance with prescribed rules, until a final configuration develops.

This chapter on self-organizing systems is restricted to Hebbian learning. The primary focus of the chapter is *principal components analysis*, which is a standard technique commonly used for data reduction in statistical pattern recognition and signal processing.

Organization of the Chapter

The material in this chapter is organized as follows. In Section 8.2 we use qualitative arguments to describe the basic principles of self-organization. This is followed by introductory material on principal components analysis in Section 8.3, which is also basic to the self-organizing systems discussed in the rest of the chapter.

With this basic background material on hand, we then proceed to study some specific self-organizing systems. In Section 8.4 we describe a simple model consisting of a single neuron, which extracts the first principal component in a self-organized manner. In Section 8.5 we describe a more elaborate self-organizing system in the form of a feedforward network with a single layer of neurons, which extracts all the principal components by building on the previous simple model. This procedure is illustrated by a computer experiment on image coding presented in Section 8.6. In Section 8.7 we describe another self-organizing system for a similar function; this system is even more elaborate because it also includes lateral connections.

In Section 8.8 we present a classification of algorithms for principal components analysis using neural networks. This is followed by Section 8.9 on the classification of data reduction algorithms into adaptive and batch methods.

In Section 8.10 we describe a nonlinear form of principal components analysis that builds on the idea of an inner-product kernel defined in accordance with Mercer's theorem, which is discussed in Chapter 6 on support vector machines.

The chapter concludes in Section 8.11 with some final thoughts on principal components analysis.

8.2 SOME INTUITIVE PRINCIPLES OF SELF-ORGANIZATION

As mentioned previously, self-organized (unsupervised) learning consists of repeatedly modifying the synaptic weights of a neural network in response to activation patterns and in accordance with prescribed rules, until a final configuration develops. The key question, of course, is how a useful configuration can finally develop from self-organization. The answer lies essentially in the following observation (Turing, 1952):

Global order can arise from local interactions.

This observation is of fundamental importance; it applies to the brain and to artificial neural networks. In particular, many originally random local interactions between neighboring neurons of a network can coalesce into states of global order and ultimately lead to coherent behavior in the form of spatial patterns or temporal rhythms; these are the essence of self-organization.

Network organization takes place at two different levels that interact with each other in the form of a *feedback* loop. The two levels are:

- *Activity*. Certain activity patterns are produced by a given network in response to input signals.
- *Connectivity*. Connection strengths (synaptic weights) of the network are modified in response to neuronal signals in the activity patterns, due to synaptic plasticity.