

Introduction to statistical pattern recognition

Overview

Statistical pattern recognition is a term used to cover all stages of an investigation from problem formulation and data collection through to discrimination and classification, assessment of results and interpretation. Some of the basic terminology is introduced and two complementary approaches to discrimination described.

1.1 Statistical pattern recognition

1.1.1 Introduction

This book describes basic pattern recognition procedures, together with practical applications of the techniques on real-world problems. A strong emphasis is placed on the statistical theory of discrimination, but clustering also receives some attention. Thus, the subject matter of this book can be summed up in a single word: ‘classification’, both supervised (using class information to design a classifier – i.e. discrimination) and unsupervised (allocating to groups without class information – i.e. clustering).

Pattern recognition as a field of study developed significantly in the 1960s. It was very much an interdisciplinary subject, covering developments in the areas of statistics, engineering, artificial intelligence, computer science, psychology and physiology, among others. Some people entered the field with a real problem to solve. The large numbers of applications, ranging from the classical ones such as automatic character recognition and medical diagnosis to the more recent ones in *data mining* (such as credit scoring, consumer sales analysis and credit card transaction analysis), have attracted considerable research effort, with many methods developed and advances made. Other researchers were motivated by the development of machines with ‘brain-like’ performance, that in some way could emulate human performance. There were many over-optimistic and unrealistic claims made, and to some extent there exist strong parallels with the

2 Introduction to statistical pattern recognition

growth of research on knowledge-based systems in the 1970s and neural networks in the 1980s.

Nevertheless, within these areas significant progress has been made, particularly where the domain overlaps with probability and statistics, and within recent years there have been many exciting new developments, both in methodology and applications. These build on the solid foundations of earlier research and take advantage of increased computational resources readily available nowadays. These developments include, for example, kernel-based methods and Bayesian computational methods.

The topics in this book could easily have been described under the term *machine learning* that describes the study of machines that can adapt to their environment and learn from example. The emphasis in machine learning is perhaps more on computationally intensive methods and less on a statistical approach, but there is strong overlap between the research areas of statistical pattern recognition and machine learning.

1.1.2 The basic model

Since many of the techniques we shall describe have been developed over a range of diverse disciplines, there is naturally a variety of sometimes contradictory terminology. We shall use the term ‘pattern’ to denote the p -dimensional data vector $\mathbf{x} = (x_1, \dots, x_p)^T$ of measurements (T denotes vector transpose), whose components x_i are measurements of the features of an object. Thus the features are the variables specified by the investigator and thought to be important for classification. In discrimination, we assume that there exist C groups or *classes*, denoted $\omega_1, \dots, \omega_C$, and associated with each pattern \mathbf{x} is a categorical variable z that denotes the class or group membership; that is, if $z = i$, then the pattern belongs to ω_i , $i \in \{1, \dots, C\}$.

Examples of patterns are measurements of an acoustic waveform in a speech recognition problem; measurements on a patient made in order to identify a disease (diagnosis); measurements on patients in order to predict the likely outcome (prognosis); measurements on weather variables (for forecasting or prediction); and a digitised image for character recognition. Therefore, we see that the term ‘pattern’, in its technical meaning, does not necessarily refer to structure within images.

The main topic in this book may be described by a number of terms such as *pattern classifier design* or *discrimination* or *allocation rule design*. By this we mean specifying the parameters of a pattern classifier, represented schematically in Figure 1.1, so that it yields the optimal (in some sense) response for a given pattern. This response is usually an estimate of the class to which the pattern belongs. We assume that we have a set of patterns of known class $\{(\mathbf{x}_i, z_i), i = 1, \dots, n\}$ (the *training* or *design* set) that we use to design the classifier (to set up its internal parameters). Once this has been done, we may estimate class membership for an unknown pattern \mathbf{x} .

The form derived for the pattern classifier depends on a number of different factors. It depends on the distribution of the training data, and the assumptions made concerning its distribution. Another important factor is the misclassification cost – the cost of making an incorrect decision. In many applications misclassification costs are hard to quantify, being combinations of several contributions such as monetary costs, time and other more subjective costs. For example, in a medical diagnosis problem, each treatment has different costs associated with it. These relate to the expense of different types of drugs,




Figure 1.1 Pattern classifier

the suffering the patient is subjected to by each course of action and the risk of further complications.

Figure 1.1 grossly oversimplifies the pattern classification procedure. Data may undergo several separate transformation stages before a final outcome is reached. These transformations (sometimes termed preprocessing, feature selection or feature extraction) operate on the data in a way that usually reduces its dimension (reduces the number of features), removing redundant or irrelevant information, and transforms it to a form more appropriate for subsequent classification. The term *intrinsic dimensionality* refers to the minimum number of variables required to capture the structure within the data. In the speech recognition example mentioned above, a preprocessing stage may be to transform the waveform to a frequency representation. This may be processed further to find formants (peaks in the spectrum). This is a feature extraction process (taking a possible nonlinear combination of the original variables to form new variables). Feature selection is the process of selecting a subset of a given set of variables.

Terminology varies between authors. Sometimes the term ‘representation pattern’ is used for the vector of measurements made on a sensor (for example, optical imager, radar) with the term ‘feature pattern’ being reserved for the small set of variables obtained by transformation (by a feature selection or feature extraction process) of the original vector of measurements. In some problems, measurements may be made directly on the feature vector itself. In these situations there is no automatic feature selection stage, with the feature selection being performed by the investigator who ‘knows’ (through experience, knowledge of previous studies and the problem domain) those variables that are important for classification. In many cases, however, it will be necessary to perform one or more transformations of the measured data.

In some pattern classifiers, each of the above stages may be present and identifiable as separate operations, while in others they may not be. Also, in some classifiers, the preliminary stages will tend to be problem-specific, as in the speech example. In this book, we consider feature selection and extraction transformations that are not application-specific. That is not to say all will be suitable for any given application, however, but application-specific preprocessing must be left to the investigator.

1.2 Stages in a pattern recognition problem

A pattern recognition investigation may consist of several stages, enumerated below. Further details are given in Appendix D. Not all stages may be present; some may be merged together so that the distinction between two operations may not be clear, even if both are carried out; also, there may be some application-specific data processing that may not be regarded as one of the stages listed. However, the points below are fairly typical.

4 Introduction to statistical pattern recognition

1. Formulation of the problem: gaining a clear understanding of the aims of the investigation and planning the remaining stages.
2. Data collection: making measurements on appropriate variables and recording details of the data collection procedure (ground truth).
3. Initial examination of the data: checking the data, calculating summary statistics and producing plots in order to get a feel for the structure.
4. Feature selection or feature extraction: selecting variables from the measured set that are appropriate for the task. These new variables may be obtained by a linear or nonlinear transformation of the original set (feature extraction). To some extent, the division of feature extraction and classification is artificial.
5. Unsupervised pattern classification or clustering. This may be viewed as exploratory data analysis and it may provide a successful conclusion to a study. On the other hand, it may be a means of preprocessing the data for a supervised classification procedure.
6. Apply discrimination or regression procedures as appropriate. The classifier is designed using a training set of exemplar patterns.
7. Assessment of results. This may involve applying the trained classifier to an independent *test set* of labelled patterns.
8. Interpretation.

The above is necessarily an iterative process: the analysis of the results may pose further hypotheses that require further data collection. Also, the cycle may be terminated at different stages: the questions posed may be answered by an initial examination of the data or it may be discovered that the data cannot answer the initial question and the problem must be reformulated.

The emphasis of this book is on techniques for performing steps 4, 5 and 6.

1.3 Issues

The main topic that we address in this book concerns classifier design: given a training set of patterns of known class, we seek to design a classifier that is optimal for the expected operating conditions (the test conditions).

There are a number of very important points to make about the sentence above, straightforward as it seems. The first is that we are given a *finite* design set. If the classifier is too complex (there are too many free parameters) it may model noise in the design set. This is an example of *over-fitting*. If the classifier is not complex enough, then it may fail to capture structure in the data. An example of this is the fitting of a set of data points by a polynomial curve. If the degree of the polynomial is too high, then, although the curve may pass through or close to the data points, thus achieving a low fitting error, the fitting curve is very variable and models every fluctuation in the data

(due to noise). If the degree of the polynomial is too low, the fitting error is large and the underlying variability of the curve is not modelled.

Thus, achieving optimal performance on the design set (in terms of minimising some error criterion perhaps) is not required: it may be possible, in a classification problem, to achieve 100% classification accuracy on the design set but the *generalisation performance* – the expected performance on data representative of the true operating conditions (equivalently, the performance on an infinite test set of which the design set is a sample) – is poorer than could be achieved by careful design. Choosing the ‘right’ model is an exercise in *model selection*.

In practice we usually do not know what is structure and what is noise in the data. Also, training a classifier (the procedure of determining its parameters) should not be considered as a separate issue from model selection, but it often is.

A second point about the design of optimal classifiers concerns the word ‘optimal’. There are several ways of measuring classifier performance, the most common being error rate, although this has severe limitations. Other measures, based on the closeness of the estimates of the probabilities of class membership to the true probabilities, may be more appropriate in many cases. However, many classifier design methods usually optimise alternative criteria since the desired ones are difficult to optimise directly. For example, a classifier may be trained by optimising a squared error measure and assessed using error rate.

Finally, we assume that the training data are representative of the test conditions. If this is not so, perhaps because the test conditions may be subject to noise not present in the training data, or there are changes in the population from which the data are drawn (population drift), then these differences must be taken into account in classifier design.

1.4 Supervised versus unsupervised

There are two main divisions of classification: *supervised classification* (or discrimination) and *unsupervised classification* (sometimes in the statistics literature simply referred to as classification or clustering).

In supervised classification we have a set of data samples (each consisting of measurements on a set of variables) with associated labels, the class types. These are used as exemplars in the classifier design.

Why do we wish to design an automatic means of classifying future data? Cannot the same method that was used to label the design set be used on the test data? In some cases this may be possible. However, even if it were possible, in practice we may wish to develop an automatic method to reduce labour-intensive procedures. In other cases, it may not be possible for a human to be part of the classification process. An example of the former is in industrial inspection. A classifier can be trained using images of components on a production line, each image labelled carefully by an operator. However, in the practical application we would wish to save a human operator from the tedious job, and hopefully make it more reliable. An example of the latter reason for performing a classification automatically is in radar target recognition of objects. For

6 Introduction to statistical pattern recognition

vehicle recognition, the data may be gathered by positioning vehicles on a turntable and making measurements from all aspect angles. In the practical application, a human may not be able to recognise an object reliably from its radar image, or the process may be carried out remotely.

In unsupervised classification, the data are not labelled and we seek to find groups in the data and the features that distinguish one group from another. Clustering techniques, described further in Chapter 10, can also be used as part of a supervised classification scheme by defining prototypes. A clustering scheme may be applied to the data for each class separately and representative samples for each group within the class (the group means, for example) used as the prototypes for that class.

1.5 Approaches to statistical pattern recognition

The problem we are addressing in this book is primarily one of pattern classification. Given a set of measurements obtained through observation and represented as a pattern vector \mathbf{x} , we wish to assign the pattern to one of C possible classes ω_i , $i = 1, \dots, C$. A *decision rule* partitions the measurement space into C regions Ω_i , $i = 1, \dots, C$. If an observation vector is in Ω_i then it is assumed to belong to class ω_i . Each region may be multiply connected – that is, it may be made up of several disjoint regions. The boundaries between the regions Ω_i are the *decision boundaries* or *decision surfaces*. Generally, it is in regions close to these boundaries that the highest proportion of misclassifications occurs. In such situations, we may reject the pattern or withhold a decision until further information is available so that a classification may be made later. This option is known as the *reject option* and therefore we have $C + 1$ outcomes of a decision rule (the reject option being denoted by ω_0) in a C -class problem.

In this section we introduce two approaches to discrimination that will be explored further in later chapters. The first assumes a knowledge of the underlying class-conditional probability density functions (the probability density function of the feature vectors for a given class). Of course, in many applications these will usually be unknown and must be estimated from a set of correctly classified samples termed the *design* or *training* set. Chapters 2 and 3 describe techniques for estimating the probability density functions explicitly.

The second approach introduced in this section develops decision rules that use the data to estimate the decision boundaries directly, without explicit calculation of the probability density functions. This approach is developed in Chapters 4, 5 and 6 where specific techniques are described.

1.5.1 Elementary decision theory

Here we introduce an approach to discrimination based on knowledge of the probability density functions of each class. Familiarity with basic probability theory is assumed. Some basic definitions are given in Appendix E.

Bayes decision rule for minimum error

Consider C classes, $\omega_1, \dots, \omega_C$, with *a priori* probabilities (the probabilities of each class occurring) $p(\omega_1), \dots, p(\omega_C)$, assumed known. If we wish to minimise the probability of making an error and we have no information regarding an object other than the class probability distribution then we would assign an object to class ω_j if

$$p(\omega_j) > p(\omega_k) \quad k = 1, \dots, C; \quad k \neq j$$

This classifies all objects as belonging to one class. For classes with equal probabilities, patterns are assigned arbitrarily between those classes.

However, we do have an *observation vector* or *measurement vector* \mathbf{x} and we wish to assign \mathbf{x} to one of the C classes. A decision rule based on probabilities is to assign \mathbf{x} to class ω_j if the probability of class ω_j given the observation \mathbf{x} , $p(\omega_j|\mathbf{x})$, is greatest over all classes $\omega_1, \dots, \omega_C$. That is, assign \mathbf{x} to class ω_j if

$$p(\omega_j|\mathbf{x}) > p(\omega_k|\mathbf{x}) \quad k = 1, \dots, C; \quad k \neq j \quad (1.1)$$

This decision rule partitions the measurement space into C regions $\Omega_1, \dots, \Omega_C$ such that if $\mathbf{x} \in \Omega_j$ then \mathbf{x} belongs to class ω_j .

The *a posteriori* probabilities $p(\omega_j|\mathbf{x})$ may be expressed in terms of the *a priori* probabilities and the class-conditional density functions $p(\mathbf{x}|\omega_i)$ using Bayes' theorem (see Appendix E) as

$$p(\omega_i|\mathbf{x}) = \frac{p(\mathbf{x}|\omega_i)p(\omega_i)}{p(\mathbf{x})}$$

and so the decision rule (1.1) may be written: assign \mathbf{x} to ω_j if

$$p(\mathbf{x}|\omega_j)p(\omega_j) > p(\mathbf{x}|\omega_k)p(\omega_k) \quad k = 1, \dots, C; \quad k \neq j \quad (1.2)$$

This is known as Bayes' rule for *minimum error*.

For two classes, the decision rule (1.2) may be written

$$l_r(\mathbf{x}) = \frac{p(\mathbf{x}|\omega_1)}{p(\mathbf{x}|\omega_2)} > \frac{p(\omega_2)}{p(\omega_1)} \text{ implies } \mathbf{x} \in \text{class } \omega_1$$

The function $l_r(\mathbf{x})$ is the *likelihood ratio*. Figures 1.2 and 1.3 give a simple illustration for a two-class discrimination problem. Class ω_1 is normally distributed with zero mean and unit variance, $p(\mathbf{x}|\omega_1) = N(\mathbf{x}|0, 1)$ (see Appendix E). Class ω_2 is a *normal mixture* (a weighted sum of normal densities) $p(\mathbf{x}|\omega_2) = 0.6N(\mathbf{x}|1, 1) + 0.4N(\mathbf{x}|-1, 2)$. Figure 1.2 plots $p(\mathbf{x}|\omega_i)p(\omega_i)$, $i = 1, 2$, where the priors are taken to be $p(\omega_1) = 0.5$, $p(\omega_2) = 0.5$. Figure 1.3 plots the likelihood ratio $l_r(\mathbf{x})$ and the threshold $p(\omega_2)/p(\omega_1)$. We see from this figure that the decision rule (1.2) leads to a disjoint region for class ω_2 .

The fact that the decision rule (1.2) minimises the error may be seen as follows. The probability of making an error, $p(\text{error})$, may be expressed as

$$p(\text{error}) = \sum_{i=1}^C p(\text{error}|\omega_i)p(\omega_i) \quad (1.3)$$

8 Introduction to statistical pattern recognition




Figure 1.2 $p(x|\omega_i)p(\omega_i)$, for classes ω_1 and ω_2




Figure 1.3 Likelihood function

where $p(\text{error}|\omega_i)$ is the probability of misclassifying patterns from class ω_i . This is given by

$$p(\text{error}|\omega_i) = \int_{\mathcal{C}[\Omega_i]} p(\mathbf{x}|\omega_i) d\mathbf{x} \quad (1.4)$$

the integral of the class-conditional density function over $\mathcal{C}[\Omega_i]$, the region of measurement space outside Ω_i (\mathcal{C} is the complement operator), i.e. $\sum_{j=1, j \neq i}^C \Omega_j$. Therefore, we

may write the probability of misclassifying a pattern as

$$\begin{aligned}
 p(\text{error}) &= \sum_{i=1}^C \int_{\mathcal{C}[\Omega_i]} p(\mathbf{x}|\omega_i) p(\omega_i) d\mathbf{x} \\
 &= \sum_{i=1}^C p(\omega_i) \left(1 - \int_{\Omega_i} p(\mathbf{x}|\omega_i) d\mathbf{x} \right) \\
 &= 1 - \sum_{i=1}^C p(\omega_i) \int_{\Omega_i} p(\mathbf{x}|\omega_i) d\mathbf{x}
 \end{aligned} \tag{1.5}$$

from which we see that minimising the probability of making an error is equivalent to maximising

$$\sum_{i=1}^C p(\omega_i) \int_{\Omega_i} p(\mathbf{x}|\omega_i) d\mathbf{x} \tag{1.6}$$

the probability of correct classification. Therefore, we wish to choose the regions Ω_i so that the integral given in (1.6) is a maximum. This is achieved by selecting Ω_i to be the region for which $p(\omega_i)p(\mathbf{x}|\omega_i)$ is the largest over all classes and the probability of correct classification, c , is

$$c = \int \max_i p(\omega_i) p(\mathbf{x}|\omega_i) d\mathbf{x} \tag{1.7}$$

where the integral is over the whole of the measurement space, and the Bayes error is

$$e_B = 1 - \int \max_i p(\omega_i) p(\mathbf{x}|\omega_i) d\mathbf{x} \tag{1.8}$$

This is illustrated in Figures 1.4 and 1.5. Figure 1.4 plots the two distributions $p(\mathbf{x}|\omega_i)$, $i = 1, 2$ (both normal with unit variance and means ± 0.5), and Figure 1.5 plots the functions $p(\mathbf{x}|\omega_i)p(\omega_i)$ where $p(\omega_1) = 0.3$, $p(\omega_2) = 0.7$. The Bayes decision boundary is marked with a vertical line at x_B . The areas of the hatched regions in Figure 1.4 represent the probability of error: by equation (1.4), the area of the horizontal hatching is the probability of classifying a pattern from class 1 as a pattern from class 2 and the area of the vertical hatching the probability of classifying a pattern from class 2 as class 1. The sum of these two areas, weighted by the priors (equation (1.5)), is the probability of making an error.

Bayes decision rule for minimum error – reject option

As we have stated above, an error or misrecognition occurs when the classifier assigns a pattern to one class when it actually belongs to another. In this section we consider the reject option. Usually it is the uncertain classifications which mainly contribute to the error rate. Therefore, rejecting a pattern (withholding a decision) may lead to a reduction in the error rate. This rejected pattern may be discarded, or set aside until further information allows a decision to be made. Although the option to reject may alleviate or remove the problem of a high misrecognition rate, some otherwise correct

10 Introduction to statistical pattern recognition




Figure 1.4 Class-conditional densities for two normal distributions




Figure 1.5 Bayes decision boundary for two normally distributed classes with unequal priors

classifications are also converted into rejects. Here we consider the trade-offs between error rate and reject rate.

Firstly, we partition the sample space into two complementary regions: R , a *reject region*, and A , an *acceptance* or *classification region*. These are defined by

$$R = \left\{ \mathbf{x} \mid 1 - \max_i p(\omega_i | \mathbf{x}) > t \right\}$$

$$A = \left\{ \mathbf{x} \mid 1 - \max_i p(\omega_i | \mathbf{x}) \leq t \right\}$$




Figure 1.6 Illustration of acceptance and reject regions

where t is a threshold. This is illustrated in Figure 1.6 using the same distributions as those in Figures 1.4 and 1.5. The smaller the value of the threshold t , the larger is the reject region R . However, if t is chosen such that

$$1 - t < \frac{1}{C}$$

or equivalently,

$$t > \frac{C - 1}{C}$$

where C is the number of classes, then the reject region is empty. This is because the minimum value which $\max_i p(\omega_i|x)$ can attain is $1/C$ (since $1 = \sum_{i=1}^C p(\omega_i|x) \leq C \max_i p(\omega_i|x)$), when all classes are equally likely. Therefore, for the reject option to be activated, we must have $t \leq (C - 1)/C$.

Thus, if a pattern \mathbf{x} lies in the region A , we classify it according to the Bayes rule for minimum error (equation (1.2)). However, if \mathbf{x} lies in the region R , we reject \mathbf{x} .

The probability of correct classification, $c(t)$, is a function of the threshold, t , and is given by equation (1.7), where now the integral is over the acceptance region, A , only

$$c(t) = \int_A \max_i [p(\omega_i)p(\mathbf{x}|\omega_i)] d\mathbf{x}$$

and the unconditional probability of rejecting a measurement \mathbf{x} , r , also a function of the threshold t , is

$$r(t) = \int_R p(\mathbf{x}) d\mathbf{x} \quad (1.9)$$

12 Introduction to statistical pattern recognition

Therefore, the error rate, e (the probability of accepting a point for classification and incorrectly classifying it), is

$$\begin{aligned} e(t) &= \int_A (1 - \max_i p(\omega_i|\mathbf{x}))p(\mathbf{x}) d\mathbf{x} \\ &= 1 - c(t) - r(t) \end{aligned}$$

Thus, the error rate and reject rate are inversely related. Chow (1970) derives a simple functional relationship between $e(t)$ and $r(t)$ which we quote here without proof. Knowing $r(t)$ over the complete range of t allows $e(t)$ to be calculated using the relationship

$$e(t) = - \int_0^t s dr(s) \quad (1.10)$$

The above result allows the error rate to be evaluated from the reject function for the Bayes optimum classifier. The reject function can be calculated using unlabelled data and a practical application is to problems where labelling of gathered data is costly.

Bayes decision rule for minimum risk

In the previous section, the decision rule selected the class for which the *a posteriori* probability, $p(\omega_j|\mathbf{x})$, was the greatest. This minimised the probability of making an error. We now consider a somewhat different rule that minimises an expected *loss* or risk. This is a very important concept since in many applications the costs associated with misclassification depend upon the true class of the pattern and the class to which it is assigned. For example, in a medical diagnosis problem in which a patient has back pain, it is far worse to classify a patient with severe spinal abnormality as healthy (or having mild back ache) than the other way round.

We make this concept more formal by introducing a loss that is a measure of the cost of making the decision that a pattern belongs to class ω_i when the true class is ω_j . We define a loss matrix Λ with components

$$\lambda_{ji} = \text{cost of assigning a pattern } \mathbf{x} \text{ to } \omega_i \text{ when } \mathbf{x} \in \omega_j$$

In practice, it may be very difficult to assign costs. In some situations, λ may be measured in monetary units that are quantifiable. However, in many situations, costs are a combination of several different factors measured in different units – money, time, quality of life. As a consequence, they may be the subjective opinion of an expert. The *conditional risk* of assigning a pattern \mathbf{x} to class ω_i is defined as

$$l^i(\mathbf{x}) = \sum_{j=1}^C \lambda_{ji} p(\omega_j|\mathbf{x})$$

The average risk over region Ω_i is

$$\begin{aligned} r^i &= \int_{\Omega_i} l^i(\mathbf{x})p(\mathbf{x}) d\mathbf{x} \\ &= \int_{\Omega_i} \sum_{j=1}^C \lambda_{ji} p(\omega_j|\mathbf{x})p(\mathbf{x}) d\mathbf{x} \end{aligned}$$

and the overall expected cost or *risk* is

$$r = \sum_{i=1}^C r^i = \sum_{i=1}^C \int_{\Omega_i} \sum_{j=1}^C \lambda_{ji} p(\omega_j | \mathbf{x}) p(\mathbf{x}) d\mathbf{x} \quad (1.11)$$

The above expression for the risk will be minimised if the regions Ω_i are chosen such that if

$$\sum_{j=1}^C \lambda_{ji} p(\omega_j | \mathbf{x}) p(\mathbf{x}) \leq \sum_{j=1}^C \lambda_{jk} p(\omega_j | \mathbf{x}) p(\mathbf{x}) \quad k = 1, \dots, C \quad (1.12)$$

then $\mathbf{x} \in \Omega_i$. This is the *Bayes decision rule for minimum risk*, with Bayes risk, r^* , given by

$$r^* = \int_{\mathbf{x}} \min_{i=1, \dots, C} \sum_{j=1}^C \lambda_{ji} p(\omega_j | \mathbf{x}) p(\mathbf{x}) d\mathbf{x}$$

One special case of the loss matrix $\mathbf{\Lambda}$ is the *equal cost* loss matrix for which

$$\lambda_{ij} = \begin{cases} 1 & i \neq j \\ 0 & i = j \end{cases}$$

Substituting into (1.12) gives the decision rule: assign \mathbf{x} to class ω_i if

$$\sum_{j=1}^C p(\omega_j | \mathbf{x}) p(\mathbf{x}) - p(\omega_i | \mathbf{x}) p(\mathbf{x}) \leq \sum_{j=1}^C p(\omega_j | \mathbf{x}) p(\mathbf{x}) - p(\omega_k | \mathbf{x}) p(\mathbf{x}) \quad k = 1, \dots, C$$

that is,

$$p(\mathbf{x} | \omega_i) p(\omega_i) \geq p(\mathbf{x} | \omega_k) p(\omega_k) \quad k = 1, \dots, C$$

implies that $\mathbf{x} \in \text{class } \omega_i$; this is the Bayes rule for minimum error.

Bayes decision rule for minimum risk – reject option

As with the Bayes rule for minimum error, we may also introduce a reject option, by which the reject region, R , is defined by

$$R = \left\{ \mathbf{x} \mid \min_i l^i(\mathbf{x}) > t \right\}$$

where t is a threshold. The decision is to accept a pattern \mathbf{x} and assign it to class ω_i if

$$l^i(\mathbf{x}) = \min_j l^j(\mathbf{x}) \leq t$$

and to reject \mathbf{x} if

$$l^i(\mathbf{x}) = \min_j l^j(\mathbf{x}) > t$$

14 Introduction to statistical pattern recognition

This decision is equivalent to defining a reject region Ω_0 with a constant conditional risk

$$l^0(\mathbf{x}) = t$$

so that the Bayes decision rule is: assign \mathbf{x} to class ω_i if

$$l^i(\mathbf{x}) \leq l^j(\mathbf{x}) \quad j = 0, 1, \dots, C$$

with Bayes risk

$$r^* = \int_R t p(\mathbf{x}) d\mathbf{x} + \int_A \min_{i=1, \dots, C} \sum_{j=1}^C \lambda_{ji} p(\omega_j | \mathbf{x}) p(\mathbf{x}) d\mathbf{x} \quad (1.13)$$

Neyman–Pearson decision rule

An alternative to the Bayes decision rules for a two-class problem is the Neyman–Pearson test. In a two-class problem there are two possible types of error that may be made in the decision process. We may classify a pattern of class ω_1 as belonging to class ω_2 or a pattern from class ω_2 as belonging to class ω_1 . Let the probability of these two errors be ϵ_1 and ϵ_2 respectively, so that

$$\epsilon_1 = \int_{\Omega_2} p(\mathbf{x} | \omega_1) d\mathbf{x} = \text{error probability of Type I}$$

and

$$\epsilon_2 = \int_{\Omega_1} p(\mathbf{x} | \omega_2) d\mathbf{x} = \text{error probability of Type II}$$

The Neyman–Pearson decision rule is to minimise the error ϵ_1 subject to ϵ_2 being equal to a constant, ϵ_0 , say.

If class ω_1 is termed the positive class and class ω_2 the negative class, then ϵ_1 is referred to as the *false negative rate*, the proportion of positive samples incorrectly assigned to the negative class; ϵ_2 is the *false positive rate*, the proportion of negative samples classed as positive.

An example of the use of the Neyman–Pearson decision rule is in radar detection where the problem is to detect a signal in the presence of noise. There are two types of error that may occur; one is to mistake noise for a signal present. This is called a *false alarm*. The second type of error occurs when a signal is actually present but the decision is made that only noise is present. This is a *missed detection*. If ω_1 denotes the signal class and ω_2 denotes the noise then ϵ_2 is the probability of false alarm and ϵ_1 is the probability of missed detection. In many radar applications, a threshold is set to give a fixed probability of false alarm and therefore the Neyman–Pearson decision rule is the one usually used.

We seek the minimum of

$$r = \int_{\Omega_2} p(\mathbf{x} | \omega_1) d\mathbf{x} + \mu \left\{ \int_{\Omega_1} p(\mathbf{x} | \omega_2) d\mathbf{x} - \epsilon_0 \right\}$$

where μ is a Lagrange multiplier¹ and ϵ_0 is the specified false alarm rate. The equation may be written

$$r = (1 - \mu\epsilon_0) + \int_{\Omega_1} \{\mu p(\mathbf{x}|\omega_2) d\mathbf{x} - p(\mathbf{x}|\omega_1) d\mathbf{x}\}$$

This will be minimised if we choose Ω_1 such that the integrand is negative, i.e.

$$\text{if } \mu p(\mathbf{x}|\omega_2) - p(\mathbf{x}|\omega_1) < 0, \quad \text{then } \mathbf{x} \in \Omega_1$$

or, in terms of the likelihood ratio,

$$\text{if } \frac{p(\mathbf{x}|\omega_1)}{p(\mathbf{x}|\omega_2)} > \mu, \quad \text{then } \mathbf{x} \in \Omega_1 \quad (1.14)$$

Thus the decision rule depends only on the within-class distributions and ignores the *a priori* probabilities.

The threshold μ is chosen so that

$$\int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} = \epsilon_0,$$

the specified false alarm rate. However, in general μ cannot be determined analytically and requires numerical calculation.

Often, the performance of the decision rule is summarised in a receiver operating characteristic (ROC) curve, which plots the true positive against the false positive (that is, the probability of detection ($1 - \epsilon_1 = \int_{\Omega_1} p(\mathbf{x}|\omega_1) d\mathbf{x}$) against the probability of false alarm ($\epsilon_2 = \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x}$) as the threshold μ is varied. This is illustrated in Figure 1.7 for the univariate case of two normally distributed classes of unit variance and means separated by a distance, d . All the ROC curves pass through the (0, 0) and (1, 1) points and as the separation increases the curve moves into the top left corner. Ideally, we would like 100% detection for a 0% false alarm rate; the closer a curve is to this the better.

For the two-class case, the minimum risk decision (see equation (1.12)) defines the decision rules on the basis of the likelihood ratio ($\lambda_{ii} = 0$):

$$\text{if } \frac{p(\mathbf{x}|\omega_1)}{p(\mathbf{x}|\omega_2)} > \frac{\lambda_{21}p(\omega_2)}{\lambda_{12}p(\omega_1)}, \quad \text{then } \mathbf{x} \in \Omega_1 \quad (1.15)$$

The threshold defined by the right-hand side will correspond to a particular point on the ROC curve that depends on the misclassification costs and the prior probabilities.

In practice, precise values for the misclassification costs will be unavailable and we shall need to assess the performance over a range of expected costs. The use of the ROC curve as a tool for comparing and assessing classifier performance is discussed in Chapter 8.

¹The method of Lagrange's undetermined multipliers can be found in most textbooks on mathematical methods, for example Wylie and Barrett (1995).




Figure 1.7 Receiver operating characteristic for two univariate normal distributions of unit variance and separation d ; $1 - \epsilon_1 = \int_{\Omega_1} p(\mathbf{x}|\omega_1) d\mathbf{x}$ is the true positive (the probability of detection) and $\epsilon_2 = \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x}$ is the false positive (the probability of false alarm)

Minimax criterion

The Bayes decision rules rely on a knowledge of both the within-class distributions and the prior class probabilities. However, situations may arise where the relative frequencies of new objects to be classified are unknown. In this situation a *minimax* procedure may be employed. The name *minimax* is used to refer to procedures for which either the maximum expected loss *or* the maximum of the error probability is a minimum. We shall limit our discussion below to the two-class problem and the minimum error probability procedure.

Consider the Bayes rule for minimum error. The decision regions Ω_1 and Ω_2 are defined by

$$p(\mathbf{x}|\omega_1)p(\omega_1) > p(\mathbf{x}|\omega_2)p(\omega_2) \text{ implies } \mathbf{x} \in \Omega_1 \quad (1.16)$$

and the Bayes minimum error, e_B , is

$$e_B = p(\omega_2) \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} + p(\omega_1) \int_{\Omega_2} p(\mathbf{x}|\omega_1) d\mathbf{x} \quad (1.17)$$

where $p(\omega_2) = 1 - p(\omega_1)$.

For *fixed* decision regions Ω_1 and Ω_2 , e_B is a linear function of $p(\omega_1)$ (we denote this function \tilde{e}_B) attaining its maximum on the region $[0, 1]$ either at $p(\omega_1) = 0$ or $p(\omega_1) = 1$. However, since the regions Ω_1 and Ω_2 are also dependent on $p(\omega_1)$ through the Bayes decision criterion (1.16), the dependency of e_B on $p(\omega_1)$ is more complex, and not necessarily monotonic.

If Ω_1 and Ω_2 are fixed (determined according to (1.16) for some specified $p(\omega_i)$), the error given by (1.17) will only be the Bayes minimum error for a particular value of $p(\omega_1)$, say p_1^* (see Figure 1.8). For other values of $p(\omega_1)$, the error given by (1.17)

must be greater than the minimum error. Therefore, the optimum curve touches the line at a tangent at p_1^* and is concave down at that point.

The minimax procedure aims to choose the partition Ω_1, Ω_2 , or equivalently the value of $p(\omega_1)$ so that the maximum error (on a test set in which the values of $p(\omega_i)$ are unknown) is minimised. For example, in the figure, if the partition were chosen to correspond to the value p_1^* of $p(\omega_1)$, then the maximum error which could occur would be a value of b if $p(\omega_1)$ were actually equal to unity. The minimax procedure aims to minimise this maximum value, i.e. minimise

$$\max\{\tilde{e}_B(0), \tilde{e}_B(1)\}$$

or minimise

$$\max \left\{ \int_{\Omega_2} p(\mathbf{x}|\omega_1) d\mathbf{x}, \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} \right\}$$

This is a minimum when

$$\int_{\Omega_2} p(\mathbf{x}|\omega_1) d\mathbf{x} = \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} \quad (1.18)$$

which is when $a = b$ in Figure 1.8 and the line $\tilde{e}_B(p(\omega_1))$ is horizontal and touches the Bayes minimum error curve at its peak value.

Therefore, we choose the regions Ω_1 and Ω_2 so that the probabilities of the two types of error are the same. The minimax solution may be criticised as being over-pessimistic since it is a Bayes solution with respect to the least favourable prior distribution. The




Figure 1.8 Minimax illustration

18 Introduction to statistical pattern recognition

strategy may also be applied to minimising the maximum risk. In this case, the risk is

$$\begin{aligned} & \int_{\Omega_1} [\lambda_{11}p(\omega_1|\mathbf{x}) + \lambda_{21}p(\omega_2|\mathbf{x})]p(\mathbf{x}) d\mathbf{x} + \int_{\Omega_2} [\lambda_{12}p(\omega_1|\mathbf{x}) + \lambda_{22}p(\omega_2|\mathbf{x})]p(\mathbf{x}) d\mathbf{x} \\ &= p(\omega_1) \left[\lambda_{11} + (\lambda_{12} - \lambda_{11}) \int_{\Omega_2} p(\mathbf{x}|\omega_1) d\mathbf{x} \right] \\ & \quad + p(\omega_2) \left[\lambda_{22} + (\lambda_{21} - \lambda_{22}) \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} \right] \end{aligned}$$

and the boundary must therefore satisfy

$$\lambda_{11} - \lambda_{22} + (\lambda_{12} - \lambda_{11}) \int_{\Omega_2} p(\mathbf{x}|\omega_1) d\mathbf{x} - (\lambda_{21} - \lambda_{22}) \int_{\Omega_1} p(\mathbf{x}|\omega_2) d\mathbf{x} = 0$$

For $\lambda_{11} = \lambda_{22}$ and $\lambda_{21} = \lambda_{12}$, this reduces to condition (1.18).

Discussion

In this section we have introduced a decision-theoretic approach to classifying patterns. This divides up the measurement space into decision regions and we have looked at various strategies for obtaining the decision boundaries. The optimum rule in the sense of minimising the error is the Bayes decision rule for minimum error. Introducing the costs of making incorrect decisions leads to the Bayes rule for minimum risk. The theory developed assumes that the *a priori* distributions and the class-conditional distributions are known. In a real-world task, this is unlikely to be so. Therefore approximations must be made based on the data available. We consider techniques for estimating distributions in Chapters 2 and 3. Two alternatives to the Bayesian decision rule have also been described, namely the Neyman–Pearson decision rule (commonly used in signal processing applications) and the minimax rule. Both require knowledge of the class-conditional probability density functions. The receiver operating characteristic curve characterises the performance of a rule over a range of thresholds of the likelihood ratio.

We have seen that the error rate plays an important part in decision-making and classifier performance assessment. Consequently, estimation of error rates is a problem of great interest in statistical pattern recognition. For given fixed decision regions, we may calculate the probability of error using (1.5). If these decision regions are chosen according to the Bayes decision rule (1.2), then the error is the *Bayes error rate* or *optimal error rate*. However, regardless of how the decision regions are chosen, the error rate may be regarded as a measure of a given decision rule's performance.

The Bayes error rate (1.5) requires complete knowledge of the class-conditional density functions. In a particular situation, these may not be known and a classifier may be designed on the basis of a training set of samples. Given this training set, we may choose to form estimates of the distributions (using some of the techniques discussed in Chapters 2 and 3) and thus, with these estimates, use the Bayes decision rule and estimate the error according to (1.5).

However, even with accurate estimates of the distributions, evaluation of the error requires an integral over a multidimensional space and may prove a formidable task. An alternative approach is to obtain bounds on the optimal error rate or distribution-free estimates. Further discussion of methods of error rate estimation is given in Chapter 8.

1.5.2 Discriminant functions

In the previous subsection, classification was achieved by applying the Bayesian decision rule. This requires knowledge of the class-conditional density functions, $p(\mathbf{x}|\omega_i)$ (such as normal distributions whose parameters are estimated from the data – see Chapter 2), or nonparametric density estimation methods (such as kernel density estimation – see Chapter 3). Here, instead of making assumptions about $p(\mathbf{x}|\omega_i)$, we make assumptions about the forms of the *discriminant functions*.

A discriminant function is a function of the pattern \mathbf{x} that leads to a classification rule. For example, in a two-class problem, a discriminant function $h(\mathbf{x})$ is a function for which

$$\begin{aligned} h(\mathbf{x}) > k &\Rightarrow \mathbf{x} \in \omega_1 \\ &< k &\Rightarrow \mathbf{x} \in \omega_2 \end{aligned} \tag{1.19}$$

for constant k . In the case of equality ($h(\mathbf{x}) = k$), the pattern \mathbf{x} may be assigned arbitrarily to one of the two classes. An optimal discriminant function for the two-class case is

$$h(\mathbf{x}) = \frac{p(\mathbf{x}|\omega_1)}{p(\mathbf{x}|\omega_2)}$$

with $k = p(\omega_2)/p(\omega_1)$. Discriminant functions are not unique. If f is a monotonic function then

$$\begin{aligned} g(\mathbf{x}) &= f(h(\mathbf{x})) > k' \Rightarrow \mathbf{x} \in \omega_1 \\ g(\mathbf{x}) &= f(h(\mathbf{x})) < k' \Rightarrow \mathbf{x} \in \omega_2 \end{aligned}$$

where $k' = f(k)$ leads to the same decision as (1.19).

In the C -group case we define C discriminant functions $g_i(\mathbf{x})$ such that

$$g_i(\mathbf{x}) > g_j(\mathbf{x}) \Rightarrow \mathbf{x} \in \omega_i \quad j = 1, \dots, C; \quad j \neq i$$

That is, a pattern is assigned to the class with the largest discriminant. Of course, for two classes, a single discriminant function

$$h(\mathbf{x}) = g_1(\mathbf{x}) - g_2(\mathbf{x})$$

with $k = 0$ reduces to the two-class case given by (1.19).

Again, we may define an optimal discriminant function as

$$g_i(\mathbf{x}) = p(\mathbf{x}|\omega_i)p(\omega_i)$$

leading to the Bayes decision rule, but as we showed for the two-class case, there are other discriminant functions that lead to the same decision.

The essential difference between the approach of the previous subsection and the discriminant function approach described here is that the form of the discriminant function is specified and is not imposed by the underlying distribution. The choice of discriminant function may depend on prior knowledge about the patterns to be classified or may be a

particular functional form whose parameters are adjusted by a training procedure. Many different forms of discriminant function have been considered in the literature, varying in complexity from the linear discriminant function (in which g is a linear combination of the x_i) to multiparameter nonlinear functions such as the multilayer perceptron.

Discrimination may also be viewed as a problem in *regression* (see Section 1.6) in which the dependent variable, y , is a class indicator and the regressors are the pattern vectors. Many discriminant function models lead to estimates of $E[y|\mathbf{x}]$, which is the aim of regression analysis (though in regression y is not necessarily a class indicator). Thus, many of the techniques we shall discuss for optimising discriminant functions apply equally well to regression problems. Indeed, as we find with feature extraction in Chapter 9 and also clustering in Chapter 10, similar techniques have been developed under different names in the pattern recognition and statistics literature.

Linear discriminant functions

First of all, let us consider the family of discriminant functions that are linear combinations of the components of $\mathbf{x} = (x_1, \dots, x_p)^T$,

$$g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0 = \sum_{i=1}^p w_i x_i + w_0 \quad (1.20)$$

This is a *linear discriminant function*, a complete specification of which is achieved by prescribing the *weight vector* \mathbf{w} and *threshold weight* w_0 . Equation (1.20) is the equation of a hyperplane with unit normal in the direction of \mathbf{w} and a perpendicular distance $|w_0|/|\mathbf{w}|$ from the origin. The value of the discriminant function for a pattern \mathbf{x} is a measure of the perpendicular distance from the hyperplane (see Figure 1.9).

A linear discriminant function can arise through assumptions of normal distributions for the class densities, with equal covariance matrices (see Chapter 2). Alternatively,




Figure 1.9 Geometry of linear discriminant function given by equation (1.20)

without making distributional assumptions, we may require the form of the discriminant function to be linear and determine its parameters (see Chapter 4).

A pattern classifier employing linear discriminant functions is termed a *linear machine* (Nilsson, 1965), an important special case of which is the *minimum-distance classifier* or nearest-neighbour rule. Suppose we are given a set of prototype points $\mathbf{p}_1, \dots, \mathbf{p}_C$, one for each of the C classes $\omega_1, \dots, \omega_C$. The minimum-distance classifier assigns a pattern \mathbf{x} to the class ω_i associated with the nearest point \mathbf{p}_i . For each point, the squared Euclidean distance is

$$|\mathbf{x} - \mathbf{p}_i|^2 = \mathbf{x}^T \mathbf{x} - 2\mathbf{x}^T \mathbf{p}_i + \mathbf{p}_i^T \mathbf{p}_i$$

and minimum-distance classification is achieved by comparing the expressions $\mathbf{x}^T \mathbf{p}_i - \frac{1}{2} \mathbf{p}_i^T \mathbf{p}_i$ and selecting the largest value. Thus, the linear discriminant function is

$$g_i(\mathbf{x}) = \mathbf{w}_i^T \mathbf{x} + w_{i0}$$

where

$$\begin{aligned} \mathbf{w}_i &= \mathbf{p}_i \\ w_{i0} &= -\frac{1}{2} |\mathbf{p}_i|^2 \end{aligned}$$

Therefore, the minimum-distance classifier is a linear machine. If the prototype points, \mathbf{p}_i , are the class means, then we have the nearest class mean classifier. Decision regions for a minimum-distance classifier are illustrated in Figure 1.10. Each boundary is the perpendicular bisector of the lines joining the prototype points of regions that are contiguous. Also, note from the figure that the decision regions are convex (that is, two arbitrary points lying in the region can be joined by a straight line that lies entirely within the region). In fact, decision regions of a linear machine are always convex. Thus, the two class problems, illustrated in Figure 1.11, although separable, cannot be separated by a linear machine. Two generalisations that overcome this difficulty are piecewise linear discriminant functions and generalised linear discriminant functions.

Piecewise linear discriminant functions

This is a generalisation of the minimum-distance classifier to the situation in which there is more than one prototype per class. Suppose there are n_i prototypes in class ω_i , $\mathbf{p}_i^1, \dots, \mathbf{p}_i^{n_i}$, $i = 1, \dots, C$. We define the discriminant function for class ω_i to be

$$g_i(\mathbf{x}) = \max_{j=1, \dots, n_i} g_i^j(\mathbf{x})$$

where g_i^j is a subsidiary discriminant function, which is linear and is given by

$$g_i^j(\mathbf{x}) = \mathbf{x}^T \mathbf{p}_i^j - \frac{1}{2} \mathbf{p}_i^{jT} \mathbf{p}_i^j \quad j = 1, \dots, n_i; i = 1, \dots, C$$

A pattern \mathbf{x} is assigned to the class for which $g_i(\mathbf{x})$ is largest; that is, to the class of the nearest prototype vector. This partitions the space into $\sum_{i=1}^C n_i$ regions known as




Figure 1.10 Decision regions for a minimum-distance classifier




Figure 1.11 Groups not separable by a linear discriminant

the Dirichlet tessellation of the space. When each pattern in the training set is taken as a prototype vector, then we have the nearest-neighbour decision rule of Chapter 3. This discriminant function generates a piecewise linear decision boundary (see Figure 1.12).

Rather than using the complete design set as prototypes, we may use a subset. Methods of reducing the number of prototype vectors (edit and condense) are described in Chapter 3, along with the nearest-neighbour algorithm. Clustering schemes may also be employed.

Generalised linear discriminant function

A *generalised linear discriminant function*, also termed a *phi machine* (Nilsson, 1965), is a discriminant function of the form

$$g(\mathbf{x}) = \mathbf{w}^T \boldsymbol{\phi} + w_0$$

where $\boldsymbol{\phi} = (\phi_1(\mathbf{x}), \dots, \phi_D(\mathbf{x}))^T$ is a vector function of \mathbf{x} . If $D = p$, the number of variables, and $\phi_i(\mathbf{x}) = x_i$, then we have a linear discriminant function.

The discriminant function is linear in the functions ϕ_i , not in the original measurements x_i . As an example, consider the two-class problem of Figure 1.13. A linear discriminant function will not separate the classes, even though they are separable. However,




Figure 1.12 Dirichlet tessellation (comprising nearest-neighbour regions for a set of prototypes) and the decision boundary (thick lines) for two classes




Figure 1.13 Nonlinear transformation of variables may permit linear discrimination

if we make the transformation

$$\begin{aligned}\phi_1(\mathbf{x}) &= x_1^2 \\ \phi_2(\mathbf{x}) &= x_2\end{aligned}$$

then the classes can be separated in the ϕ -space by a straight line. Similarly, disjoint classes can be transformed into a ϕ -space in which a linear discriminant function could separate the classes (provided that they are separable in the original space).

The problem, therefore, is simple. Make a good choice for the functions $\phi_i(\mathbf{x})$, then use a linear discriminant function to separate the classes. But, how do we choose ϕ_i ? Specific examples are shown in Table 1.1.

Clearly there is a problem in that as the number of functions that are used as a basis set increases, so does the number of parameters that must be determined using the limited

Table 1.1 Discriminant functions, ϕ

Discriminant function	Mathematical form, $\phi_i(\mathbf{x})$
linear	$\phi_i(\mathbf{x}) = x_i, i = 1, \dots, p$
quadratic	$\phi_i(\mathbf{x}) = x_{k_1}^{l_1} x_{k_2}^{l_2}, i = 1, \dots, (p+1)(p+2)/2 - 1$ $l_1, l_2 = 0 \text{ or } 1; k_1, k_2 = 1, \dots, p, l_1, l_2 \text{ not both zero}$
v th-order polynomial	$\phi_i(\mathbf{x}) = x_{k_1}^{l_1} \dots x_{k_v}^{l_v}, i = 1, \dots, \binom{p+v}{v} - 1$ $l_1, \dots, l_v = 0 \text{ or } 1; k_1, \dots, k_v = 1, \dots, p$ $l_i \text{ not all zero}$
radial basis function	$\phi_i(\mathbf{x}) = \phi(\mathbf{x} - \mathbf{v}_i)$ for centre \mathbf{v}_i and function ϕ
multilayer perceptron	$\phi_i(\mathbf{x}) = f(\mathbf{x}^T \mathbf{v}_i + v_{i0})$ for direction \mathbf{v}_i and offset v_{i0} . f is the logistic function, $f(z) = 1/(1 + \exp(-z))$

training set. A complete quadratic discriminant function requires $D = (p+1)(p+2)/2$ terms and so for C classes there are $C(p+1)(p+2)/2$ parameters to estimate. We may need to apply a constraint or ‘regularise’ the model to ensure that there is no over-fitting.

An alternative to having a set of different functions is to have a set of functions of the same parametric form, but which differ in the values of the parameters they take,

$$\phi_i(\mathbf{x}) = \phi(\mathbf{x}; \mathbf{v}_i)$$

where \mathbf{v}_i is a set of parameters. Different models arise depending on the way the variable \mathbf{x} and the parameters \mathbf{v} are combined. If

$$\phi(\mathbf{x}; \mathbf{v}) = \phi(|\mathbf{x} - \mathbf{v}|)$$

that is, ϕ is a function only of the magnitude of the difference between the pattern \mathbf{x} and the weight vector \mathbf{v} , then the resulting discriminant function is known as a *radial basis function*. On the other hand, if ϕ is a function of the scalar product of the two vectors

$$\phi(\mathbf{x}; \mathbf{v}) = \phi(\mathbf{x}^T \mathbf{v} + v_0)$$

then the discriminant function is known as a *multilayer perceptron*. It is also a model known as projection pursuit. Both the radial basis function and the multilayer perceptron models can be used in regression.

In these latter examples, the discriminant function is no longer linear in the parameters. Specific forms for ϕ for radial basis functions and for the multilayer perceptron models will be given in Chapters 5 and 6.

Summary

In a multiclass problem, a pattern \mathbf{x} is assigned to the class for which the discriminant function is the largest. A linear discriminant function divides the feature space by a

hyperplane whose orientation is determined by the weight vector \mathbf{w} and distance from the origin by the weight threshold w_0 . The decision regions produced by linear discriminant functions are convex.

A piecewise linear discriminant function permits non-convex and disjoint decision regions. Special cases are the nearest-neighbour and nearest class mean classifier.

A generalised linear discriminant function, with fixed functions ϕ_i , is linear in its parameters. It permits non-convex and multiply connected decision regions (for suitable choices of ϕ_i). Radial basis functions and multilayer perceptrons can be regarded as generalised linear discriminant functions with flexible functions ϕ_i whose parameters must be determined or specified using the training set.

The Bayes decision rule is optimal (in the sense of minimising classification error) and with sufficient flexibility in our discriminant functions we ought to be able to achieve optimal performance in principle. However, we are limited by a finite number of training samples and also, once we start to consider parametric forms for the ϕ_i , we lose the simplicity and ease of computation of the linear functions.

1.6 Multiple regression

Many of the techniques and procedures described within this book are also relevant to problems in *regression*, the process of investigating the relationship between a dependent (or response) variable Y and independent (or predictor) variables X_1, \dots, X_p ; a regression function expresses the expected value of Y in terms of X_1, \dots, X_p and model parameters. Regression is an important part of statistical pattern recognition and, although the emphasis of the book is on discrimination, practical illustrations are sometimes given on problems of a regression nature.

The discrimination problem itself is one in which we are attempting to predict the values of one variable (the class variable) given measurements made on a set of independent variables (the pattern vector, \mathbf{x}). In this case, the response variable is categorical. Posing the discrimination problem as one in regression is discussed in Chapter 4.

Regression analysis is concerned with predicting the mean value of the response variable given measurements on the predictor variables and assumes a model of the form

$$E[y|\mathbf{x}] \triangleq \int yp(y|\mathbf{x}) dy = f(\mathbf{x}; \boldsymbol{\theta})$$

where f is a (possibly nonlinear) function of the measurements \mathbf{x} and $\boldsymbol{\theta}$, a set of parameters of f . For example,

$$f(\mathbf{x}; \boldsymbol{\theta}) = \theta_0 + \boldsymbol{\theta}^T \mathbf{x}$$

where $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T$, is a model that is linear in the parameters and the variables. The model

$$f(\mathbf{x}; \boldsymbol{\theta}) = \theta_0 + \boldsymbol{\theta}^T \boldsymbol{\phi}(\mathbf{x})$$

where $\boldsymbol{\theta} = (\theta_1, \dots, \theta_D)^T$ and $\boldsymbol{\phi} = (\phi_1(\mathbf{x}), \dots, \phi_D(\mathbf{x}))^T$ is a vector of nonlinear functions of \mathbf{x} , is linear in the parameters but nonlinear in the variables. *Linear regression*




Figure 1.14 Population regression line (solid line) with representation of spread of conditional distribution (dotted lines) for normally distributed error terms, with variance depending on x

refers to a regression model that is linear in the parameters, but not necessarily in the variables.

Figure 1.14 shows a regression summary for some hypothetical data. For each value of x , there is a population of y values that varies with x . The solid line connecting the conditional means, $E[y|x]$, is the *regression line*. The dotted lines either side represent the spread of the conditional distribution (± 1 standard deviation from the mean).

It is assumed that the difference (commonly referred to as an error or residual), ϵ_i , between the measurement on the response variable and its predicted value conditional on the measurements on the predictors,

$$\epsilon_i = y_i - E[y|x_i]$$

is an unobservable random variable. A normal model for the errors (see Appendix E) is often assumed,

$$p(\epsilon) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2} \frac{\epsilon^2}{\sigma^2}\right)$$

That is,

$$p(y_i|x_i, \theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2\sigma^2}(y_i - f(x_i; \theta))^2\right)$$

Given a set of data $\{(y_i, x_i), i = 1, \dots, n\}$, the maximum likelihood estimate of the model parameters (the value of the parameters for which the data are ‘most likely’),

discussed further in Appendix B), θ , is that for which

$$p(\{(y_i, \mathbf{x}_i)\}|\theta)$$

is a maximum. Assuming independent samples, this amounts to determining the value of θ for which the commonly used least squares error,

$$\sum_{i=1}^n (y_i - f(\mathbf{x}_i; \theta))^2 \quad (1.21)$$

is a minimum (see the exercises at the end of the chapter).

For the linear model, procedures for estimating the parameters are described in Chapter 4.

1.7 Outline of book

The aim in writing this volume is to provide a comprehensive account of statistical pattern recognition techniques with emphasis on methods and algorithms for discrimination and classification. In recent years there have been many developments in multivariate analysis techniques, particularly in nonparametric methods for discrimination and classification. These are described in this book as extensions to the basic methodology developed over the years.

This chapter has presented some basic approaches to statistical pattern recognition. Supplementary material on probability theory and data analysis can be found in the appendices.

A road map to the book is given in Figure 1.15, which describes the basic pattern recognition cycle. The numbers in the figure refer to chapters and appendices of this book.

Chapters 2 and 3 describe basic approaches to supervised classification via Bayes' rule and estimation of the class-conditional densities. Chapter 2 considers normal-based models. Chapter 3 addresses nonparametric approaches to density estimation.

Chapters 4–7 take a discriminant function approach to supervised classification. Chapter 4 describes algorithms for linear discriminant functions. Chapter 5 considers kernel-based approaches for constructing nonlinear discriminant functions, namely radial basis functions and support vector machine methods. Chapter 6 describes alternative, projection-based methods, including the multilayer perceptron neural network. Chapter 7 describes tree-based approaches.

Chapter 8 addresses the important topic of performance assessment: how good is your designed classifier and how well does it compare with competing techniques? Can improvement be achieved with an ensemble of classifiers?

Chapters 9 and 10 consider techniques that may form part of an exploratory data analysis. Chapter 9 describes methods of feature selection and extraction, both linear and nonlinear. Chapter 10 addresses *unsupervised* classification or clustering.

Finally, Chapter 11 covers additional topics on pattern recognition including model selection.

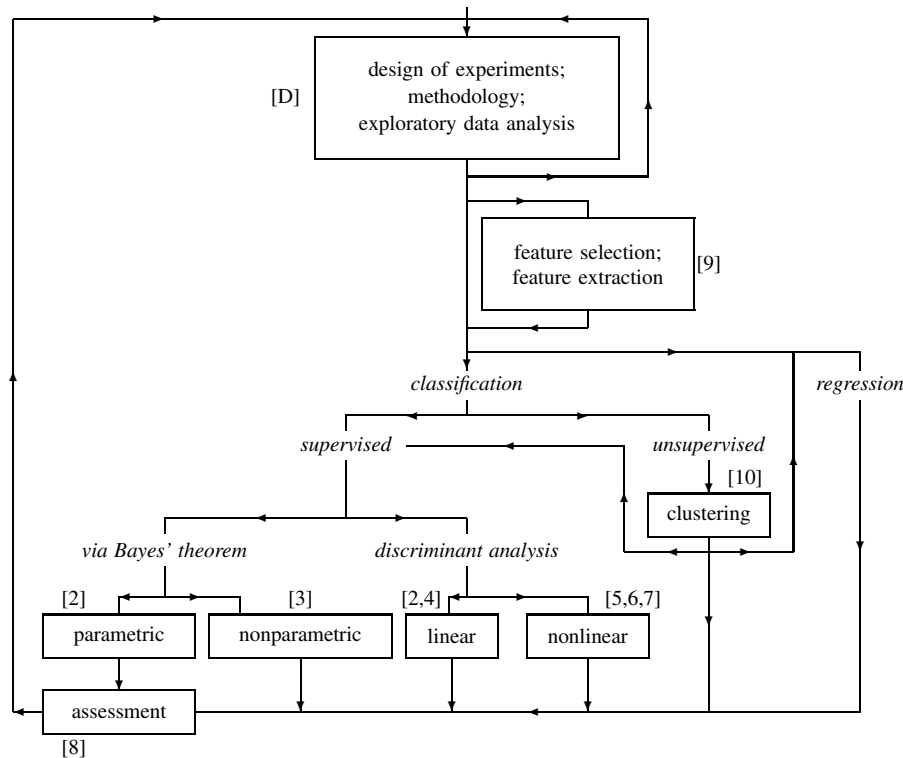


Figure 1.15 The pattern recognition cycle; numbers in parentheses refer to chapters and appendices of this book

1.8 Notes and references

There was a growth of interest in techniques for automatic pattern recognition in the 1960s. Many books appeared in the early 1970s, some of which are still very relevant today and have been revised and reissued. More recently, there has been another flurry of books on pattern recognition, particularly incorporating developments in neural network methods.

A very good introduction is provided by the book of Hand (1981a). Perhaps a little out of date now, it provides nevertheless a very readable account of techniques for discrimination and classification written from a statistical point of view and is to be recommended. Two of the main textbooks on statistical pattern recognition are those by Fukunaga (1990) and Devijver and Kittler (1982). Written perhaps with an engineering emphasis, Fukunaga's book provides a comprehensive account of the most important aspects of pattern recognition, with many examples, computer projects and problems. Devijver and Kittler's book covers the nearest-neighbour decision rule and feature selection and extraction in some detail, though not at the neglect of other important areas of statistical pattern recognition. It contains detailed mathematical accounts of techniques and algorithms, treating some areas in depth.

Another important textbook is that by Duda *et al.* (2001). Recently revised, this presents a thorough account of the main topics in pattern recognition, covering many recent developments. Other books that are an important source of reference material are those by Young and Calvert (1974), Tou and Gonzales (1974) and Chen (1973). Also, good accounts are given by Andrews (1972), a more mathematical treatment, and Therrien (1989), an undergraduate text.

Recently, there have been several books that describe the developments in pattern recognition that have taken place over the last decade, particularly the 'neural network' aspects, relating these to the more traditional methods. A comprehensive treatment of neural networks is provided by Haykin (1994). Bishop (1995) provides an excellent introduction to neural network methods from a statistical pattern recognition perspective. Ripley's (1996) account provides a thorough description of pattern recognition from within a statistical framework. It includes neural network methods, approaches developed in the field of machine learning, recent advances in statistical techniques as well as development of more traditional pattern recognition methods and gives valuable insights into many techniques gained from practical experience. Hastie *et al.* (2001) provide a thorough description of modern techniques in pattern recognition. Other books that deserve a mention are those by Schalkoff (1992) and Pao (1989).

Hand (1997) gives a short introduction to pattern recognition techniques and the central ideas in discrimination and places emphasis on the comparison and assessment of classifiers.

A more specialised treatment of discriminant analysis and pattern recognition is the book by McLachlan (1992a). This is a very good book. It is not an introductory textbook, but provides a thorough account of recent advances and sophisticated developments in discriminant analysis. Written from a statistical perspective, the book is a valuable guide to theoretical and practical work on statistical pattern recognition and is to be recommended for researchers in the field.

Comparative treatments of pattern recognition techniques (statistical, neural and machine learning methods) are provided in the volume edited by Michie *et al.* (1994) who report on the outcome of the *Statlog* project. Technical descriptions of the methods are given, together with the results of applying those techniques to a wide range of problems. This volume provides the most extensive comparative study available. More than 20 different classification procedures were considered for about 20 data sets.

The book by Watanabe (1985), unlike the books above, is not an account of statistical methods of discrimination, though some are included. Rather, it considers a wider perspective of human cognition and learning. There are many other books in this latter area. Indeed, in the early days of pattern recognition, many of the meetings and conferences covered the humanistic and biological side of pattern recognition in addition to the mechanical aspects. Although these non-mechanical aspects are beyond the scope of this book, the monograph by Watanabe provides one unifying treatment that we recommend for background reading.

There are many other books on pattern recognition. Some of those treating more specific parts (such as clustering) are cited in the appropriate chapter of this book. In addition, most textbooks on multivariate analysis devote some attention to discrimination and classification. These provide a valuable source of reference and are cited elsewhere in the book.

The website www.statistical-pattern-recognition.net contains references and links to further information on techniques and applications.

Exercises

In some of the exercises, it will be necessary to generate samples from a multivariate density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Many computer packages offer routines for this. However, it is a simple matter to generate samples from a normal distribution with unit variance and zero mean (for example, Press *et al.*, 1992). Given a vector \mathbf{Y}_i of such samples, then the vector $\mathbf{U}\boldsymbol{\Lambda}^{1/2}\mathbf{Y}_i + \boldsymbol{\mu}$ has the required distribution, where \mathbf{U} is the matrix of eigenvectors of the covariance matrix and $\boldsymbol{\Lambda}^{1/2}$ is the diagonal matrix whose diagonal elements are the square roots of the corresponding eigenvalues (see Appendix C).

1. Consider two multivariate normally distributed classes,

$$p(\mathbf{x}|\omega_i) = \frac{1}{(2\pi)^{p/2}|\boldsymbol{\Sigma}_i|^{1/2}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^T \boldsymbol{\Sigma}_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i) \right\}$$

with means $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ and equal covariance matrices, $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}$. Show that the logarithm of the likelihood ratio is linear in the feature vector \mathbf{x} . What is the equation of the decision boundary?

2. Determine the equation of the decision boundary for the more general case of $\boldsymbol{\Sigma}_1 = \alpha\boldsymbol{\Sigma}_2$, for scalar α (normally distributed classes as in Exercise 1). In particular, for two univariate distributions, $N(0, 1)$ and $N(1, 1/4)$, show that one of the decision regions is bounded and determine its extent.
3. For the distributions in Exercise 1, determine the equation of the minimum risk decision boundary for the loss matrix

$$\boldsymbol{\Lambda} = \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$$

4. Consider two multivariate normally distributed classes (ω_2 with mean $(-1, 0)^T$ and ω_1 with mean $(1, 0)^T$, and identity covariance matrix). For a given threshold μ (see equation (1.14)) on the likelihood ratio, determine the regions Ω_1 and Ω_2 in a Neyman–Pearson rule.
5. Consider three bivariate normal distributions, $\omega_1, \omega_2, \omega_3$ with identity covariance matrices and means $(-2, 0)^T, (0, 0)^T$ and $(0, 2)^T$. Show that the decision boundaries are piecewise linear. Now define a class A as the *mixture* of ω_1 and ω_3 ,

$$p_A(\mathbf{x}) = 0.5p(\mathbf{x}|\omega_1) + 0.5p(\mathbf{x}|\omega_3)$$

and class B as bivariate normal with identity covariance matrix and mean $(a, b)^T$, for some a, b . What is the equation of the Bayes decision boundary? Under what conditions is it piecewise linear?

6. Consider two uniform distributions with equal priors

$$p(x|\omega_1) = \begin{cases} 1 & \text{when } 0 \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

$$p(x|\omega_2) = \begin{cases} \frac{1}{2} & \text{when } \frac{1}{2} \leq x \leq \frac{5}{2} \\ 0 & \text{otherwise} \end{cases}$$

Show that the reject function is given by

$$r(t) = \begin{cases} \frac{3}{8} & \text{when } 0 \leq t \leq \frac{1}{3} \\ 0 & \text{when } \frac{1}{3} < t \leq 1 \end{cases}$$

Hence calculate the error rate by integrating (1.10).

7. Reject option. Consider two classes, each normally distributed with means $x = 1$ and $x = -1$ and unit variances; $p(\omega_1) = p(\omega_2) = 0.5$. Generate a test set and use it (without using class labels) to estimate the reject rate as a function of the threshold t . Hence, estimate the error rate for no rejection. Compare with the estimate based on a labelled version of the test set. Comment on the use of this procedure when the true distributions are unknown and the densities have to be estimated.
8. The area of a sphere of radius r in p dimensions, S_p , is

$$S_p = \frac{2\pi^{\frac{p}{2}} r^{p-1}}{\Gamma(p/2)}$$

where Γ is the gamma function ($\Gamma(1/2) = \pi^{1/2}$, $\Gamma(1) = 1$, $\Gamma(x+1) = x\Gamma(x)$). Show that the probability of a sample, \mathbf{x} , drawn from a zero-mean normal distribution with covariance matrix $\sigma^2 \mathbf{I}$ (\mathbf{I} is the identity matrix) and having $|\mathbf{x}| \leq R$ is

$$\int_0^R S_p(r) \frac{1}{(2\pi\sigma^2)^{p/2}} \exp\left(-\frac{r^2}{2\sigma^2}\right) dr$$

Evaluate this numerically for $R = 2\sigma$ and for $p = 1, \dots, 10$. What do the results tell you about the distribution of normal samples in high-dimensional spaces?

9. In a two-class problem, let the cost of misclassifying a class ω_1 pattern be C_1 and the cost of misclassifying a class ω_2 pattern be C_2 . Show that the point on the ROC curve that minimises the risk has gradient

$$\frac{C_2 p(\omega_2)}{C_1 p(\omega_1)}$$

10. Show that under the assumption of normally distributed residuals, the maximum likelihood solution for the parameters of a linear model is equivalent to minimising the sum-square error (1.21).

