

## ON ERROR-RATE ESTIMATION IN NONPARAMETRIC CLASSIFICATION

Anil K. Ghosh and Peter Hall

*Indian Statistical Institute, Kolkata and University of Melbourne*

*Abstract:* There is a substantial literature on the estimation of error rate, or risk, for nonparametric classifiers. Error-rate estimation has at least two purposes: accurately describing the error rate, and estimating the tuning parameters that permit the error rate to be minimised. In the light of work on related problems in nonparametric statistics, it is attractive to argue that both problems admit the same solution. Indeed, methods for optimising the point-estimation performance of nonparametric curve estimators often start from an accurate estimator of error. However, we argue in this paper that accurate estimators of error rate in classification tend to give poor results when used to choose tuning parameters; and vice versa. Concise theory is used to illustrate this point in the case of cross-validation (which gives very accurate estimators of error rate, but poor estimators of tuning parameters) and the smoothed bootstrap (where error-rate estimation is poor but tuning-parameter approximations are particularly good). The theory is readily extended to other methods, for example to the 0.632+ bootstrap approach, which gives good estimators of error rate but poor estimators of tuning parameters. Reasons for the apparent contradiction are given, and numerical results are used to point to the practical implications of the theory.

*Key words and phrases:* Bayes risk, bootstrap, cross-validation, discrimination, error rate, kernel methods, nonparametric density estimation, risk.

### 1. Introduction

Cross-validation is a widely used technique for estimating the risk, or error rate, of classification procedures. It often gives estimators which are close to unbiased, and which have good mean squared error properties. In fact, cross-validation is frequently viewed as the method to beat when alternative approaches to risk estimation are suggested. See, for example, Efron (1983) and Efron and Tibshirani (1997).

One might expect that good performance in estimating risk would be accompanied by good performance in determining the values of tuning parameters that minimise risk. Indeed, in a number of related model-selection problems, computing a good estimator of error is the first step in approximating the values of the parameters that minimise error. Early work of this type includes that of Hall (1983), Bowman (1984), Stone (1984) and Faraway and Jhun (1990).

However, in important ways the problem of risk estimation in classification is significantly different from a number of apparently similar problems in nonparametric statistics. A major reason is that the tuning parameters used to construct classifiers may influence performance only in a relatively small number of places, for example the places where population densities cross. Therefore, the impact that the parameters have on risk can be relatively minor. As a result, in classification problems, it is possible to construct a particularly accurate estimator of risk which is of very little value for estimating tuning parameters; cross-validation turns out to be of this type. The contrary case also arises — empirical risk-based methods for estimating tuning parameters may perform that task very well, but give poor estimators of risk.

It is unsurprising that this problem is not well understood. Indeed, it is somewhat contradictory to argue that one should not seek an accurate estimator of risk when attempting to minimise that quantity empirically. However, a consequence of not fully understanding the problem is that methods such as cross-validation, and its jackknife or bootstrap competitors, which are designed to minimise risk, are in practice pressed into service to select tuning parameters. This can be inappropriate.

In the present paper we point to the shortcomings of cross-validation for estimating tuning parameters for classifiers, and also to the advantages of certain other approaches that give accurate estimators of tuning parameters but poor estimators of risk. In order to make our discussion and technical arguments transparent, we treat a relatively simple, univariate problem, where standard kernel estimators are used as the basis for classifiers. However, similar results can be derived in multivariate settings, and also when methods other than kernel estimators are used for classification. In the kernel case, the tuning parameters referred to above are bandwidths.

In this paper we derive theoretical results that address the following points.

- (a) The cross-validation estimator of risk is root- $n$  consistent and, in fact, is asymptotically equivalent to a nonparametric maximum likelihood estimator.
- (b) Notwithstanding property (a), the part of the cross-validation estimator that depends on the tuning parameters is very highly stochastically variable, so much so that it has an unboundedly large number of local extrema which bear no important asymptotic relationship to the parameters that actually minimise risk.
- (c) In marked contrast to (a), a smoothed bootstrap estimator of risk is highly biased, and in consequence can have poor convergence rates relative to the cross-validation approach.
- (d) Despite the drawbacks noted in (c), the smoothed bootstrap method produces accurate estimators of the parameters that minimise risk.
- (e) The smoothed bootstrap method is particularly robust against inappropriate choice of smoothing parameters, and in fact those quantities can be selected

within a very broad range without appreciably influencing performance. Technical proofs of our results are given in Ghosh and Hall (2006).

Property (e) will be reinforced by our numerical work in Section 3, which will also introduce an adaptive, empirical approach to smoothing the bootstrap. Property (b) has been discussed by Hall and Kang (2005), but without a concise mathematical account of the issues involved. The erratic way in which the cross-validation criterion varies with tuning parameters is well known.

More generally, the theoretical results given in the present paper can be augmented by others, which show that some of the problems that afflict cross-validation can be reduced by applying Breiman's (1996) bagging technique to dampen the effects of excessive variability. However, it seems difficult to achieve good performance of the smoothed bootstrap approach, and for that reason we omit from this paper a theoretical treatment of bagged cross-validation.

The fact that bagging does not redress all the problems associated with cross-validation is highlighted in our numerical work, where we treat three additional approaches: bagged cross-validation, and the bootstrap methods suggested by Efron (1983) and Efron and Tibshirani (1997), respectively. In terms of their performance at estimating tuning parameters, the first and second of these techniques lie between cross-validation (at the lower-performance end of the scale) and the smoothed bootstrap (at the upper end). The method of Efron and Tibshirani (1997) gives particularly poor estimators of tuning parameters.

There is an especially large literature on nonparametric classification. Kernel-based approaches date from work of Fix and Hodges (1951). More generally, a large variety of classification methodologies has been developed based on empirical forms of the Bayes classifier. Relatively recent contributions include those of Chanda and Ruymgaart (1989), Krzyżak (1991), Lapko (1993), Pawlak (1993), Lugosi and Pawlak (1994), Devroye, Györfi and Lugosi (1996), Lugosi and Nobel (1996), Ancukiewicz (1998), Yang (1999a,b), Mammen and Tsybakov (1999), Steele and Patterson (2000), and Lin (2001).

## 2. Main Results

### 2.1. Error rates and their estimators

Let  $F$  and  $G$  denote distributions with respective densities  $f$  and  $g$ , and let  $\mathcal{X} = \{X_1, \dots, X_m\}$  and  $\mathcal{Y} = \{Y_1, \dots, Y_n\}$  be datasets drawn respectively from  $F$  and  $G$ . We consider ways of classifying a new data value,  $x$  say, to either the  $F$  or the  $G$  population.

Assume that the distributions  $F$  and  $G$  have respective prior probabilities  $p$  and  $1 - p$ . Define

$$\Delta = p f - (1 - p) g. \quad (2.1)$$

The Bayes classifier,  $\mathcal{A}_0$  say, allocates  $x$  to  $F$  or  $G$  according as  $\Delta(x)$  is positive or negative, respectively. The corresponding error rate, or risk, for classification of data on a compact interval  $\mathcal{I}$ , is

$$\text{err}_{\mathcal{A}_0} = p \int_{\mathcal{I}} I\{\Delta(x) < 0\} f(x) dx + (1-p) \int_{\mathcal{I}} I\{\Delta(x) > 0\} g(x) dx.$$

A general class of classifiers can be constructed by replacing  $\Delta$  by an estimator,  $\hat{\Delta}$ . As a prelude to defining  $\hat{\Delta}$  we introduce density estimators  $\hat{f}$  and  $\hat{g}$ , and their leave-one-out versions  $\hat{f}_{-i}$  and  $\hat{g}_{-i}$ . Let  $K$  be a nonnegative kernel and  $h_1$  and  $h_2$  be bandwidths, and put  $m_1 = m - 1$ ,  $n_1 = n - 1$ ,

$$\begin{aligned} \hat{f}(x) &= \frac{1}{mh_1} \sum_{i=1}^m K\left(\frac{x - X_i}{h_1}\right), & \hat{g}(x) &= \frac{1}{nh_2} \sum_{i=1}^n K\left(\frac{x - Y_i}{h_2}\right), \\ \hat{f}_{-i}(x) &= \frac{1}{m_1 h_1} \sum_{j=1}^{(i)} K\left(\frac{x - X_j}{h_1}\right), & \hat{g}_{-i}(x) &= \frac{1}{n_1 h_2} \sum_{j=1}^{(i)} K\left(\frac{x - Y_j}{h_2}\right), \end{aligned} \quad (2.2)$$

where  $\sum_j^{(i)}$  denotes summation over indices  $j$  not equal to  $i$ . Define

$$\hat{\Delta} = p \hat{f} - (1-p) \hat{g}.$$

The empirical classifier  $\mathcal{A}_1$ , which assigns  $x$  to distribution  $F$  if  $\hat{\Delta}(x) > 0$ , and to  $G$  otherwise, has the following empirical risk:

$$\text{emperr}_{\mathcal{A}_1}(h_1, h_2) = p \int_{\mathcal{I}} I\{\hat{\Delta}(x) < 0\} f(x) dx + (1-p) \int_{\mathcal{I}} I\{\hat{\Delta}(x) > 0\} g(x) dx.$$

Its average value over all possible datasets  $\mathcal{X}$  and  $\mathcal{Y}$ , i.e. its expected value, is

$$\text{err}_{\mathcal{A}_1}(h_1, h_2) = p \int_{\mathcal{I}} P\{\hat{\Delta}(x) < 0\} f(x) dx + (1-p) \int_{\mathcal{I}} P\{\hat{\Delta}(x) > 0\} g(x) dx.$$

The cross-validation estimator of  $\text{err}_{\mathcal{A}_1}$  is

$$\begin{aligned} \text{CV}(h_1, h_2) &= \frac{p}{m} \sum_{i=1}^m I\{\hat{\Delta}_{f,-i}(X_i) < 0, X_i \in \mathcal{I}\} \\ &\quad + \frac{1-p}{n} \sum_{i=1}^n I\{\hat{\Delta}_{g,-i}(Y_i) > 0, Y_i \in \mathcal{I}\}, \end{aligned} \quad (2.3)$$

where  $\hat{\Delta}_{f,-i} = p \hat{f}_{-i} - (1-p) \hat{g}$  and  $\hat{\Delta}_{g,-i} = p \hat{f} - (1-p) \hat{g}_{-i}$ .

An alternative way of estimating risk is to resample from smoothed versions of the empirical distributions of  $X$  and  $Y$ , conditional on the data in  $\mathcal{Z} = \mathcal{X} \cup \mathcal{Y}$ , using bootstrap methods. Specifically, let  $\tilde{f}$  and  $\tilde{g}$  denote the versions of  $\hat{f}$  and

$\hat{g}$ , respectively, at (2.2), constructed using bandwidths  $h_3$  and  $h_4$  in place of  $h_1$  and  $h_2$ . (Potentially, the bandwidths employed at this point can be quite different from those used to construct the classifier.) Let  $\mathcal{X}^* = \{X_1^*, \dots, X_m^*\}$  and  $\mathcal{Y}^* = \{Y_1^*, \dots, Y_n^*\}$  denote datasets drawn by sampling randomly, conditional on  $\mathcal{Z}$ , from the distributions with respective densities  $\tilde{f}$  and  $\tilde{g}$ . Construct the versions  $\hat{f}^*$  and  $\hat{g}^*$  of  $\hat{f}$  and  $\hat{g}$  from these resamples, on this occasion using the original bandwidths  $h_1$  and  $h_2$ :

$$\hat{f}^*(x) = \frac{1}{mh_1} \sum_{i=1}^m K\left(\frac{x - X_i^*}{h_1}\right), \quad \hat{g}^*(x) = \frac{1}{nh_2} \sum_{i=1}^n K\left(\frac{x - Y_i^*}{h_2}\right).$$

Put  $\hat{\Delta}^* = p\hat{f}^* - (1-p)\hat{g}^*$ . The bootstrap classifier consists of assigning a new data value  $x$  to  $F$  if  $\hat{\Delta}^*(x) > 0$ , and assigning it to  $G$  otherwise.

The long-run error rate of this classifier, conditional on the data  $\mathcal{Z}$ , is given by

$$\widehat{\text{err}}_{\mathcal{A}_1}(h_1, h_2) = p \int_{\mathcal{I}} P\{\hat{\Delta}^*(x) < 0 \mid \mathcal{Z}\} \tilde{f}(x) dx + (1-p) \int_{\mathcal{I}} P\{\hat{\Delta}^*(x) > 0 \mid \mathcal{Z}\} \tilde{g}(x) dx.$$

In particular,  $\widehat{\text{err}}_{\mathcal{A}_1}$  is an approximation to  $\text{err}_{\mathcal{A}_1}$  alternative to the cross-validation criterion, CV.

Results related to those described in subsequent sections can be derived for discriminative nonparametric approaches to classification, rather than the generative one discussed above. For example, denoting the  $X$  and  $Y$  populations by 0 and 1, respectively, taking  $W = 0$  or 1 in this context, and writing  $Z$  to denote a new data value (either an  $X$  or a  $Y$ ), we can treat the problem of estimating  $\pi(z) \equiv P(W = 0 \mid Z = z)$  as one of regression rather than density estimation. One option would be to run a local-polynomial regression smoother through the data pairs  $(0, X_1), \dots, (0, X_m), (1, Y_1), \dots, (1, Y_n)$ , where each value is interpreted as an observation of  $(W, Z)$  (here,  $p$  is taken to equal  $m/(m+n)$ ), and use the estimated regression function,  $\hat{\pi}$ , to effect classification by assigning a new datum  $Z$  to  $X$  if and only if  $\hat{\pi}(Z) > 1/2$ . The properties that we discuss below, relating (for example) to the high degree of variability of cross-validation for choosing bandwidth, all have parallels in the setting of this discriminative method.

## 2.2. Summary of properties of CV, emperr and $\widehat{\text{err}}$ as approximations to risk

It is known (see e.g., Hall and Kang (2005)) that the bandwidths that are optimal in the sense of minimising the risk  $\text{err}_{\mathcal{A}_1}$  of the empirical classification rule  $\mathcal{A}_1$ , are generally of the same size as  $h = n^{-1/5}$ . Therefore we assess risks when the bandwidths are on this scale.

In this context, the “regret” of the rule  $\mathcal{A}_1$  is of size  $h^4$ . That is, when both  $h_1$  and  $h_2$  are of size  $n^{-1/5}$ , the difference between the risk  $\text{err}_{\mathcal{A}_1}$  for  $\mathcal{A}_1$ , and the risk  $\text{err}_{\mathcal{A}_0}$  for the Bayes classifier, is asymptotic to a constant multiple of  $h^4$ . (See e.g., Hall and Kang (2005), and also (2.14) below.) The regret captures all of the influence that bandwidth choice has on risk. Moreover the empirical risk,  $\text{emperr}_{\mathcal{A}_1}$ , also equals the Bayes risk plus a term of order  $h^4$ . However, in the case of  $\text{emperr}_{\mathcal{A}_1}$ , the  $h^4$  term includes non-negligible stochastic fluctuations that depend substantially on  $h_1$  and  $h_2$ , but in a manner that does not reflect the dependence of  $\text{err}_{\mathcal{A}_1}$  on those bandwidths.

We show that the cross-validation estimator of risk equals a term which does not depend on either of the bandwidths  $h_1$  and  $h_2$ , plus a highly stochastically volatile quantity which is of size  $h^{7/2}$ , and so is an order of magnitude larger than  $h^4$ . More particularly, the fluctuations of the cross-validation criterion  $\text{CV}$ , as a function of the bandwidths  $h_1$  and  $h_2$ , are largely unrelated to the ways in which either the risk  $\text{err}_{\mathcal{A}_1}$ , or its empirical form  $\text{emperr}_{\mathcal{A}_1}$ , are influenced by the bandwidths. Therefore, minimising the cross-validation criterion does not correspond, in an asymptotic or in another meaningful sense, to minimising either the true or the empirical risk.

However, we note that  $\text{CV}(h_1, h_2)$  is close to being an unbiased estimator of  $\text{err}_{\mathcal{A}_1}(h_1, h_2)$ , indeed so close that it captures all the main effects of bandwidth choice on risk. Therefore, the difficulties that afflict cross-validation arise from stochastic variability, not systematic error.

The strengths and weaknesses of  $\widehat{\text{err}}_{\mathcal{A}_1}$ , as an alternative to  $\text{CV}$ , are diametrically opposite to those of  $\text{CV}$ . In particular,  $\widehat{\text{err}}_{\mathcal{A}_1}$  suffers from substantial bias as an estimator  $\text{err}_{\mathcal{A}_1}$ , but the stochastic variability of that portion of  $\widehat{\text{err}}_{\mathcal{A}_1}$  that captures the main effects of bandwidth is particularly low. As a result,  $\widehat{\text{err}}_{\mathcal{A}_1}$  can be used to effectively choose bandwidths for density-based classifiers.

### 2.3. Details of properties of cross-validation

We assume that:

$K$  is symmetric, compactly supported, integrates to 1 and has two bounded derivatives; the function  $\Delta$ , defined at (2.1), vanishes in  $\mathcal{I}$  only at  $r$  isolated points, say  $y_1, \dots, y_r$ , in the interior of  $\mathcal{I}$ , and at each point  $y_i$ ,  $\Delta'(y_i) \neq 0$ ;  $f$  and  $g$  are continuous on the real line, and have two Hölder-continuous derivatives in an open neighbourhood of  $y_i$  for  $1 \leq i \leq r$ ;  $m$  and  $n$  increase together, and  $n/m \rightarrow \rho$ , where  $0 < \rho < \infty$ . (2.4)

Put  $h = n^{-1/5}$ . Given  $B > 1$ , let  $\mathcal{H} = \mathcal{H}(B)$  denote the set of values of  $hu$  for which  $B^{-1} \leq u \leq B$ . We take both the bandwidths  $h_1$  and  $h_2$ , used to construct the estimators at (2.2), to be in  $\mathcal{H}$ .

For  $1 \leq i \leq r$ , let  $\mathcal{W}_{X_i}$ ,  $\mathcal{W}_{Y_i}$ ,  $\mathcal{G}_{X_i}$  and  $\mathcal{G}_{Y_i}$  denote independent stochastic processes, with  $\mathcal{W}_{X_i}$  and  $\mathcal{W}_{Y_i}$  being standard Wiener processes, and  $\mathcal{G}_{X_i}$  and  $\mathcal{G}_{Y_i}$  Gaussian processes having zero means and covariance given by, for both  $Z = X$  and  $Z = Y$ ,

$$\text{Cov} \{ \mathcal{G}_{Z_i}(t^{(1)}), \mathcal{G}_{Z_i}(t^{(2)}) \} = \int K(st^{(1)}) K(st^{(2)}) ds. \tag{2.5}$$

Put  $\kappa = \int K^2$  and  $\kappa_2 = \int x^2 K(x) dx$ . Given  $0 < \rho < \infty$  and  $u_1, u_2 > 0$ , let  $u = (u_1, u_2)$ ,

$$d(y | u) = \frac{1}{2} \kappa_2 \{ p u_1^2 f''(y) - (1 - p) u_2^2 g''(y) \}, \tag{2.6}$$

$$\mathcal{V}_i(u) = \frac{1}{\Delta'(y_i)} \left[ p \{ \rho f(y_i) \} \frac{1}{2} \mathcal{G}_{X_i}(u_1) + (1 - p) g(y_i) \frac{1}{2} \mathcal{G}_{Y_i}(u_2) + d(y_i | u) \right], \tag{2.7}$$

$$T(u) = \sum_{i=1}^r \left[ p \{ \rho f(y_i) \} \frac{1}{2} \mathcal{W}_{X_i} \{ \mathcal{V}_i(u) \} + (1 - p) g(y_i) \frac{1}{2} \mathcal{W}_{Y_i} \{ \mathcal{V}_i(u) \} \right],$$

$$\begin{aligned} \tau(u) &= \frac{1}{2} \kappa \sum_{i=1}^r |\Delta'(y_i)|^{-1} \{ p^2 \rho f(y_i) u_1^{-1} + (1 - p)^2 g(y_i) u_2^{-1} \} \\ &\quad + \frac{1}{2} \sum_{i=1}^r |\Delta'(y_i)|^{-1} d(y_i | u)^2. \end{aligned} \tag{2.8}$$

If  $\mathcal{I} = [a, b]$ , define  $y_0 = a$  and  $y_{r+1} = b$  and, for  $1 \leq i \leq r + 1$ , let  $L_i$  denote the number of indices  $j$  for which  $y_{i-1} < X_j < y_i$  if  $\Delta < 0$  on  $(y_{i-1}, y_i)$ , or the number of  $j$  for which  $y_{i-1} < Y_j < y_i$  if  $\Delta > 0$  on  $(y_{i-1}, y_i)$ . Put  $p_i = p/m$  in the first of these cases, and  $p_i = (1 - p)/n$  in the second.

**Theorem 2.1.** *Assume conditions (2.4). Then the stochastic processes  $\mathcal{W}_{X_i}$ ,  $\mathcal{W}_{Y_i}$  and  $\mathcal{V}_i$  can be constructed, depending in each instance on  $n$ , such that, with  $h_j = hu_j$  for  $j = 1$  and  $2$ ,*

$$\text{CV}(h_1, h_2) = \sum_{i=1}^{r+1} p_i L_i + h^{\frac{7}{2}} T(u) + o_p(h^{\frac{7}{2}}), \tag{2.9}$$

where the remainder is of the stated order uniformly in  $B^{-1} \leq u_1, u_2 \leq B$ , with  $B > 1$ . Furthermore,

$$E\{\text{CV}(h_1, h_2)\} = \text{err}_{\mathcal{A}_1} + o(h^4) \tag{2.10}$$

$$= \text{err}_{\mathcal{A}_0} + \frac{1}{2} \sum_{i=1}^r |\Delta'(y_i)|^{-1} E\{ p \hat{f}(y_i) - (1 - p) \hat{g}(y_i) \}^2 + o(h^4) \tag{2.11}$$

$$= \text{err}_{\mathcal{A}_0} + h^4 \tau(u) + o(h^4), \tag{2.12}$$

again uniformly in  $B^{-1} \leq u_1, u_2 \leq B$ .

The two main terms on the right-hand side of (2.9) represent a division of  $CV(h_1, h_2)$  into parts that represent, respectively, the dominant part of CV that does not depend on the bandwidths  $h_1$  and  $h_2$ , and the dominant part that is influenced by those bandwidths. In particular, the series  $\sum_i p_i L_i$  does not depend on  $h_1$  and  $h_2$ . Its expected value equals the risk of the Bayes rule for classification on the interval  $\mathcal{I}$ :

$$E\left(\sum_{i=1}^{r+1} p_i L_i\right) = \text{err}_{\mathcal{A}_0}.$$

The variance of  $\sum_i p_i L_i$  is of order  $n^{-1}$ .

If we were given the values of  $y_1, \dots, y_r$ , and told also the signs of  $\Delta$  on the intervals between adjacent values of  $y_i$ , then the nonparametric maximum likelihood estimator of  $\text{err}_{\mathcal{A}_0}$  would be exactly  $\sum_i p_i L_i$ . In particular, this series has minimum variance among all unbiased estimators of  $\text{err}_{\mathcal{A}_0}$ , and its convergence rate,  $n^{-1/2}$ , cannot be improved upon. Therefore, the series  $\sum_i p_i L_i$ , which represents the dominant part of CV and converges to its expected value at the slower of the rates for the two respective terms in (2.9), cannot be made significantly more accurate as an approximation to the Bayes-rule risk  $\text{err}_{\mathcal{A}_0}$ .

However, the term  $\sum_i p_i L_i$  does not provide any information about the effect of bandwidth on classification performance in neighbourhoods of the points  $y_1, \dots, y_r$ . Of course, that information is crucial to understanding how properties of the classifier are influenced by its construction. We have to pass to the term  $h^{7/2} T(u)$  on the right-hand side of (2.9) in order to obtain any information about how  $h_1$  and  $h_2$  influence  $CV(h_1, h_2)$ .

Revealingly, the second term varies stochastically in a very erratic manner. Indeed, since the Wiener processes  $\mathcal{W}_{X_i}$  and  $\mathcal{W}_{Y_i}$  have fractal sample paths then, with probability 1,  $T(u_1, u_2)$  has an infinite number of local minima, as a function of  $u_1$  and  $u_2$ , in any rectangle. Indicative of this high degree of volatility, a graph of a realisation of the function  $v = T(u_1, u_2)$ , as a function of  $u_1$  and  $u_2$ , is, with probability 1, a surface of fractal dimension exceeding 2.

These difficulties persist even if we take  $u_1 = u_2 = t$ , say. (A choice that reflects the common practice of using the same bandwidth to construct density estimators from either dataset.) Of course, in practice the cross-validation criterion  $CV(ht, ht)$  has only a finite number of local maxima in any nondegenerate interval, but the fact that the stochastic approximant  $T(t, t)$  has an infinite number of local maxima there implies that the number of local minima of  $CV(ht, ht)$  in the interval increases without bound as sample size diverges.

These results provide a theoretical explanation of the observed high degree of stochastic variability of the cross-validation criterion, in terms of the way



it describes the effect of bandwidth choice on risk. The results also indicate why choosing the bandwidths to minimise CV is fraught with practical difficulty. When using cross-validation with real data it is found that the criterion has many local minima, few of which seem more appropriate than the others, and that this problem becomes more, rather than less, pronounced as sample size increases. Both these properties are implied by the theoretical results discussed in the two previous paragraphs, and so the theory provides insight into practice.

On the other hand, results (2.10)–(2.12) assert that the cross-validation criterion is close to being an unbiased approximation to the risk of the empirical rule based on  $\hat{\Delta}$ .

#### 2.4. Discussion of (2.4) and its multivariate version

If it is not true that, as prescribed by (2.4),  $\Delta$  “vanishes in  $\mathcal{I}$  only at  $r$  isolated points...  $y_1, \dots, y_r$  in the interior of  $\mathcal{I}$ , and [each]  $\Delta'(y_i) \neq 0$ ” (call this assumption (A)), then the most natural alternative condition would arguably be that “ $\Delta \equiv 0$  on a nondegenerate subinterval,  $\mathcal{I}'$  say, of  $\mathcal{I}$ .” In this case the classification problem becomes substantially more difficult, since data that take values in  $\mathcal{I}'$  have equal likelihood of coming from either population. Reflecting this fact, the minimum size of classification error is an order of magnitude larger than  $n^{-4/5}$  (its value when (A) holds).

The  $d$ -variate analogue of (A) is the following: the set,  $\mathcal{J}$  say, of points  $y$  for which  $\Delta(y) = 0$ , is of codimension 1 and of finite measure in  $d - 1$  dimensions; the slopes at which the functions  $pf$  and  $(1 - p)g$  meet are bounded away from zero along  $\mathcal{J}$ , i.e.  $|p\{\partial f(y)/\partial y_j\} - (1 - p)\{\partial g(y)/\partial y_j\}|$  is bounded away from zero uniformly in  $1 \leq j \leq d$  and  $y \in \mathcal{J}$ . Call this assumption (B), and take the two  $d$ -variate bandwidth matrices, one for each of the populations, to be diagonal, with each component a fixed multiple of a common value  $h$ . Then it can be shown that, under assumption (B) and the condition that both densities have two derivatives, the optimal order of  $h$  for classification is  $n^{-1/(d+4)}$ . In fact, an expansion of risk has, as its dominant terms, two quantities of respective sizes  $(nh^d)^{-1}$  and  $h^4$ , leading to multivariate versions of (2.12) above and (2.13)–(2.15) below, identical to those formulae except that now  $h = n^{-1/(d+4)}$ , and  $\tau$  and  $\hat{\tau}$  have different expressions. Background to these properties is given in Sections 5.3 and 5.4 of Hall and Kang (2005), and properties of cross-validation can be derived similarly.

#### 2.5. Properties of empirical risk

Here we state and discuss a version of Theorem 2.1 for  $\text{emperr}_{\mathcal{A}_1}$ , rather than for the cross-validation approximation to the risk. The empirical risk is not computable in practice, but from some viewpoints one would not expect the

cross-validation approximation to be markedly inferior to the empirical risk, at least in terms of the way it reflects properties of the bandwidths. In fact, it is substantially inferior.

Recall that  $u = (u_1, u_2)$ , and define  $\hat{\tau} = \tau + \hat{\tau}_1$ , where  $\tau$  is as at (2.8) and

$$\hat{\tau}_1(u) = \sum_{i=1}^r \Delta'(y_i) \int \left[ I\{v < \mathcal{V}_i(u)\} - P\{v < \mathcal{V}_i(u)\} \right] dv.$$

**Theorem 2.2.** *Assume conditions (2.4). Then, with  $\mathcal{V}_i$  as in Theorem 2.1, and with  $h_j = hu_j$  for  $j = 1$  and  $2$ ,*

$$\begin{aligned} \text{emperr}_{\mathcal{A}_1}(h_1, h_2) &= \text{err}_{\mathcal{A}_1}(h_1, h_2) + h^4 \hat{\tau}_1(u) + o_p(h^4) \\ &= \text{err}_{\mathcal{A}_0} + h^4 \hat{\tau}(u) + o_p(h^4), \end{aligned} \quad (2.13)$$

where the remainders are of the stated orders uniformly in  $B^{-1} \leq u_1, u_2 \leq B$ , for any  $B > 1$ .

A formula for the error of the classification rule  $\mathcal{A}_1$  is obtainable directly from (2.10) and (2.12):

$$\text{err}_{\mathcal{A}_1} = \text{err}_{\mathcal{A}_0} + h^4 \tau(u) + o(h^4), \quad (2.14)$$

where the positive function  $\tau$  is given at (2.8). In particular, (2.14) implies that the regret is asymptotic to a constant multiple of  $h^4$ . Result (2.13) shows that the difference between the empirical risk and the actual risk is of the same size as, but not asymptotically equal to, the regret. The former, multiplied by  $h^{-4}$ , converges in distribution to a nondegenerate random variable which can take both positive and negative values, whereas the regret, multiplied by  $h^{-4}$ , converges to a positive constant.

More importantly, a comparison of (2.9), (2.13) and (2.14) shows that there is no useful connection between the parts of  $\text{CV}(h_1, h_2)$  that depend on  $h_1$  and  $h_2$ , and the corresponding parts of either the risk  $\text{err}_{\mathcal{A}_1}$  or its empirical version  $\text{emperr}_{\mathcal{A}_1}$ . In particular, the term  $h^{7/2}T(u)$  in (2.9) is an order of magnitude larger than both  $h^4 \tau(u)$  and  $h^4 \hat{\tau}(u)$ , on the right-hand sides of (2.14) and (2.13), respectively. Moreover, the fluctuations of  $T(u)$ , as a function of  $u$ , bear no relationship to those of  $\tau(u)$  or  $\hat{\tau}(u)$ . Therefore, cross-validation cannot be used effectively to choose the bandwidths that minimise either  $\text{err}_{\mathcal{A}_1}$  or  $\text{emperr}_{\mathcal{A}_1}$ .

## 2.6. Properties of bootstrap estimator of risk

In the light of what we have learned in earlier sections, the properties of  $\widehat{\text{err}}_{\mathcal{A}_1}(h_1, h_2)$  are relatively transparent. In particular, provided the bandwidths  $h_3$  and  $h_4$  (used to construct the density estimators  $\tilde{f}$  and  $\tilde{g}$ ) are sufficiently large

to ensure that  $\tilde{f}''$  and  $\tilde{g}''$  are consistent for  $f''$  and  $g''$ , respectively, the bootstrap analogue of (2.14) holds:

$$\widehat{\text{err}}_{\mathcal{A}_1}(h_1, h_2) = \widehat{\text{err}}_{\mathcal{A}_0} + h^4 \tau(u) + o_p(h^4), \quad (2.15)$$

uniformly in  $B^{-1} \leq u_1, u_2 \leq B$ . In (2.15),  $\widehat{\text{err}}_{\mathcal{A}_0}$  denotes the estimator of  $\text{err}_{\mathcal{A}_0}$  that is obtained on replacing, in the definition of  $\text{err}_{\mathcal{A}_0}$ , the unknown densities  $f$  and  $g$  by their estimators  $\tilde{f}$  and  $\tilde{g}$ :

$$\widehat{\text{err}}_{\mathcal{A}_0} = p \int_{\mathcal{I}} I\{\tilde{\Delta}(x) < 0\} \tilde{f}(x) dx + (1-p) \int_{\mathcal{I}} I\{\tilde{\Delta}(x) > 0\} \tilde{g}(x) dx,$$

where  $\tilde{\Delta} = p\tilde{f} - (1-p)\tilde{g}$ . Note particularly that  $\widehat{\text{err}}_{\mathcal{A}_0}$  does not depend on the bandwidths  $h_1$  and  $h_2$ .

The quantity  $\widehat{\text{err}}_{\mathcal{A}_0}$  will generally not be a good estimator of  $\text{err}_{\mathcal{A}_0}$ . In particular, the relatively large values needed for the bandwidths  $h_3$  and  $h_4$  will ensure that  $\widehat{\text{err}}_{\mathcal{A}_0}$  suffers from significant bias, although (e.g., under the conditions of Theorem 2.3 below) it will be consistent. However, this inaccuracy is not necessarily a problem if our aim is determine, from  $\widehat{\text{err}}_{\mathcal{A}_1}$ , the influence that  $h_1$  and  $h_2$  have on the true risk,  $\text{err}_{\mathcal{A}_1}$ . Since  $\widehat{\text{err}}_{\mathcal{A}_0}$  does not depend on  $h_1$  and  $h_2$ , then the main effect of the influence of those quantities is expressed through the term  $h^4 \tau(u)$  on the right-hand side of (2.15), and so is exactly the same as main effect of the influence of  $h_1$  and  $h_2$  on  $\text{err}_{\mathcal{A}_1}$ ; see (2.14). Hence, we can use  $\widehat{\text{err}}_{\mathcal{A}_1}$  effectively to choose the bandwidth that minimises risk.

**Theorem 2.3.** *Assume conditions (2.4), and that the bandwidths  $h_3$  and  $h_4$  both satisfy  $n^{(1/5)-\epsilon}h_j \rightarrow \infty$  and  $n^\epsilon h_j \rightarrow 0$  for some  $\epsilon > 0$ . Then (2.15) holds, uniformly in  $B^{-1} \leq u_1, u_2 \leq B$ .*

### 3. Numerical Properties

In this section we report the results of a simulation study addressing numerical properties of risk estimators based on cross-validation and the bootstrap. We know from our theoretical work that having an estimator of risk that is good for estimating bandwidth, is not necessarily the same as having a good estimator of risk itself. For example, we showed in Section 2 that the bootstrap gives an estimator of risk that is seriously biased, relative to the estimator produced by cross-validation, and that the bootstrap approach nevertheless gives better bandwidth estimators. Both these results are reflected starkly in numerical experiments, although we report here only the results about bandwidth choice.

Likewise, Efron's (1983) bootstrap method is known not to be a good estimator of risk, since it usually leads to overestimation (see e.g., Efron and Tibshirani (1997)), but here we show that it nevertheless performs reasonably well as a bandwidth selector. On the other hand, Efron and Tibshirani's (1997) 0.632+

bootstrap method gives very good estimators of risk, as that paper shows, but it does not perform well when used to select bandwidth, as we demonstrate below.

For the most part we consider the case where the distributions with densities  $f$  and  $g$  are equal-probability mixtures of two univariate normal distributions, with means 0 and 2 (in the case of  $f$ ) and 1 and 3 (for  $g$ ), each component having variance 0.25. The context where the Normal distributions are both replaced by lognormal distributions, or by Cauchy distributions, will also be discussed. We take  $p = 1/2$  and  $m = n$ , and  $h_2 = sh_1$ , where  $s$  denotes the ratio of a measure of the scale of  $f$  to that for  $g$ . This approach is often used in practice. Throughout,  $K$  is the standard Gaussian kernel.

The first panel of Figure 3.1 depicts 100 plots of  $CV(h_1, sh_1)$  as a function of  $h_1$ , when  $s$  is taken equal to 1. Since the two population distributions have the same variance then this choice of  $s$  is reasonable. We discuss shortly the case where  $s$  is estimated from data.

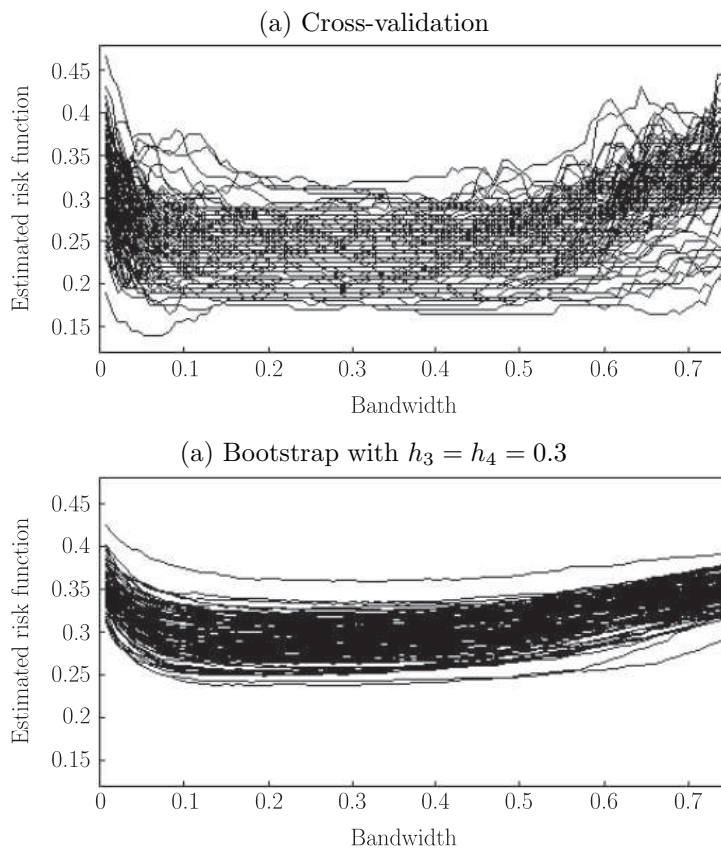


Figure 3.1. Estimated risk using (a) cross-validation method or (b) the bootstrap.

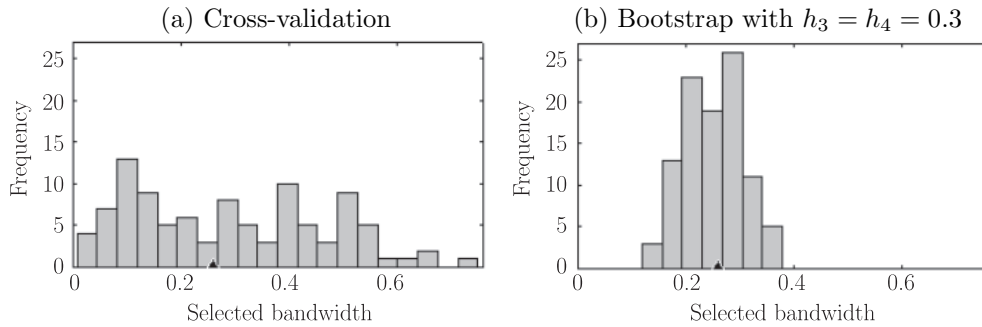


Figure 3.2. Histogram estimators of empirical bandwidth distributions, when bandwidths are selected using (a) cross-validation or (b) the bootstrap.

Each curve in Figure 3.1 is computed for a different pair  $(\mathcal{X}, \mathcal{Y})$  of random samples, of sizes  $m = n = 100$ . The second panel of the figure shows, for the same 100 sample pairs, plots of the bootstrap alternative to the cross-validation criterion,  $\widehat{\text{err}}_{\mathcal{A}_1}(h_1, sh_1)$ . We used  $h_3 = h_4 = 0.3$  when computing the bootstrap density estimators  $\hat{f}^*$  and  $\hat{g}^*$ .

The erratic nature of the graphs in the first panel of Figure 3.1 reflects the high degree of variability of cross-validation, demonstrated theoretically in Section 2. This suggests that cross-validation has substantial difficulty, relative to the bootstrap, approximating the optimal bandwidth. That is confirmed by Figure 3.2, which gives a histogram estimator of the distribution of the bandwidths that minimise  $\text{CV}(h_1, sh_1)$  (in the left-hand panel of Figure 3.2) or  $\widehat{\text{err}}_{\mathcal{A}_1}(h_1, sh_1)$  (in the right-hand panel). In this setting the theoretically optimal bandwidth, in the sense of minimising risk under the constraint  $h_2 = sh_1$ , can be shown to be  $h_1 = 0.26$ , which value is indicated by a small black triangle on the horizontal axes. The bootstrap bandwidth estimator is close to being unbiased, and has low stochastic variability, whereas the cross-validation estimator is skewed to the right and is very highly variable. Results for different sample sizes, and for other densities  $f$  and  $g$ , are similar.

In order to improve the performance of cross-validation we used the bootstrap aggregation, or bagging, technique suggested by Breiman (1996). From each sample, a proportion  $\alpha$  (where  $0 < \alpha < 1$ ) of data was resampled, without replacement, to form a new subsample. Cross-validation was applied to this subsample to estimate the risk function. This step was repeated many times, and an overall estimator of risk was obtained by averaging. The first panel of Figure 3.3 shows that the resulting bagged version of  $\text{CV}(h_1, sh_1)$  has substantially lower stochastic variation than its unbaggged counterpart.

Of course, when using the bagged form of  $\text{CV}(h_1, sh_1)$  to select bandwidth, we need to correct for the fact that we reduced sample size by the factor  $\alpha$ . As

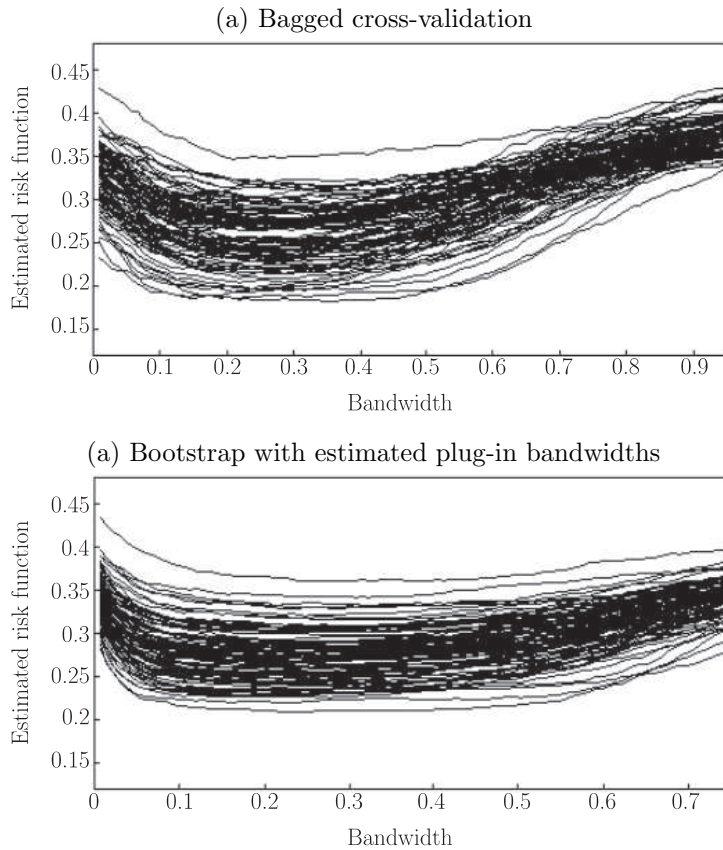


Figure 3.3. Estimated risk for (a) bagged version of cross-validation or (b) bootstrap using estimated plug-in bandwidths.

discussed in Section 2, it is known that the optimal bandwidth is of size  $n^{-1/5}$ , and so an appropriate correction is readily obtained by taking the bandwidth that minimises the bagged form of  $CV(h_1, sh_1)$ , and reducing it by multiplying by the factor  $\alpha^{1/5}$ .

As an aid to determining the appropriate value of  $\alpha$  we experimented with different training-sample sizes ( $m = n = 50, 100, 150$  and  $200$ ) and different values of  $\alpha$ . In each case, we generated 100 different training samples and computed the true risk function corresponding to the selected bandwidth. Average values of these risk functions are reported in Figure 3.4, for different values of  $\alpha$  and  $n$ . It can be seen that, for  $n \geq 100$ , the method is largely unaffected by different choices of  $\alpha$ , although values in the range  $0.2 \leq \alpha \leq 0.4$  are mildly preferable. Similar results are obtained for other density pairs  $(f, g)$ . Therefore we take  $\alpha = 0.3$  in the work below.




Figure 3.4. Average risk for the bagged cross-validation rule, expressed as a function of  $\alpha$ , equal to the proportion of the sample that is resampled.

Analogous experiments were conducted to determine appropriate formulae for the plug-in bandwidths  $h_3$  and  $h_4$  used in the bootstrap algorithm. Echoing the wide theoretical range permitted in Theorem 2.3, and reflecting our experience when determining the amount of bagging that should be applied, the effect of choice of smoothing parameter was found to be particularly small for  $n \geq 100$ . Indeed, choice of  $h_3$  and  $h_4$  does not seem to be a significant issue. Our numerical experiments suggest that if the training-sample sizes are approximately equal then the choices  $h_3 = n^{0.05} \hat{h}_1$  and  $h_4 = n^{0.05} s \hat{h}_1$  are appropriate in the bootstrap stage, where  $\hat{h}_1$  is the optimum bandwidth estimated by the bagged version of cross-validation. This approach was used to produce the second panel of Figure 3.3.

Figure 3.5 depicts the relative performances of two of the bandwidth selectors discussed above: bagged cross-validation with sampling fraction equal to 0.3, and the bootstrap with empirical choice of  $h_3$  and  $h_4$ . (Figure 3.2 gave the analogous histogram in the case of cross-validation.) For comparison we also include the bootstrap method suggested by Efron (1983), and the 0.632+ bootstrap method proposed by Efron and Tibshirani (1997).

Like standard cross-validation based on  $CV(h_1, sh_1)$ , the 0.632+ bootstrap does a good job estimating risk for its own sake and, in particular, produces estimators that are significantly less biased than those given by the bootstrap criterion  $\widehat{\text{err}}_{\mathcal{A}_1}(h_1, sh_1)$ . However, also like cross-validation, it has poor performance when used to estimate bandwidth. Likewise, Efron's (1983) method is superior to  $\widehat{\text{err}}_{\mathcal{A}_1}$  at estimating risk; the fact that it is inferior to  $\widehat{\text{err}}_{\mathcal{A}_1}$  when used to choose bandwidth is not a contradiction.

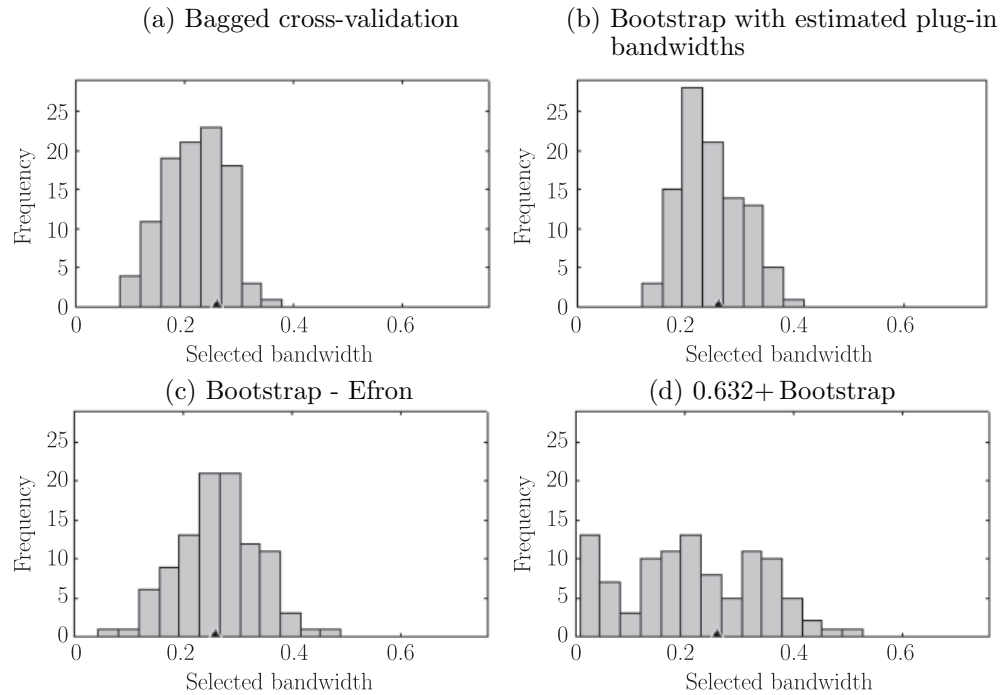


Figure 3.5. Histogram estimators of empirical bandwidth distributions, when bandwidths are selected using (a) bagged cross-validation, (b) bootstrap employing estimated plug-in bandwidths, (c) Efron's (1983) method, or (d) the 0.632+ bootstrap.

In practice one would use an estimator,  $\hat{s}$  say, of  $s$ , for example the ratio of standard deviations or of the interquartile ranges. In the normal-mixture case, results obtained for either of these approaches were virtually identical to that when  $s = 1$ . In particular, Figures 3.1–3.5 were almost unchanged. The results reported for the remainder of this section will be for the case where  $s$  was replaced by the ratio of interquartile ranges.

In practice, cross-validation leads to an estimator of error rate with multiple local minima. Ghosh and Chaudhuri (2004) suggested choosing, in such cases, the maximum of the optimisers. In the present paper we examined performance of both the minimum and the maximum of the optimisers, and again chose the latter because it gave reduced misclassification rate.

An alternative to leave-one-out cross-validation is  $\ell$ -fold cross-validation, where the data are divided randomly into  $\ell$  parts. This can be done many times, using different random partitions, and the error rates aggregated. The resulting “bagged  $\ell$ -fold cross-validation” method can be used as the basis for bandwidth selection, by minimising the bagged criterion and then multiplying



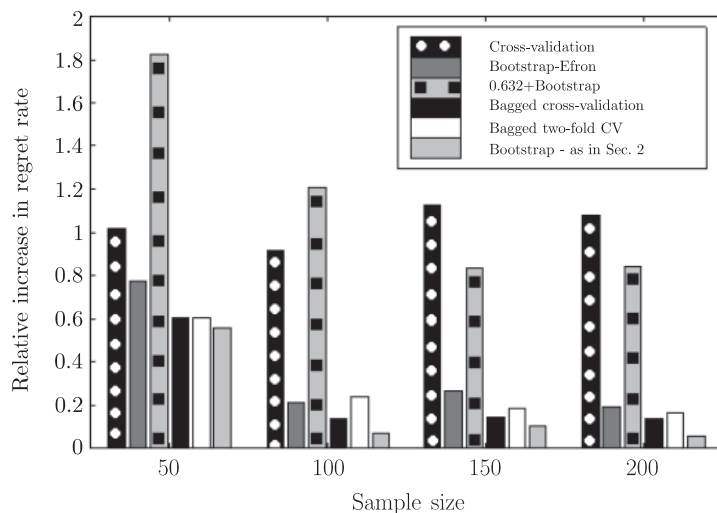


Figure 3.6. Relative increase in regret for different classifiers, in the case of normal mixture distributions.

the resulting bandwidth by the correction factor  $(1 - \ell^{-1})^{1/5}$ . We explored this approach numerically for  $\ell = 2, 5$  and  $10$ , and, as expected, found that the results obtained were similar to those for bagged leave-one-out cross-validation. (Likewise, unbaggged  $\ell$ -fold cross-validation gave results similar to those for unbaggged leave-one-out cross-validation.) We report here only the results for  $\ell = 2$ ; they are slightly better than their counterparts for either of the other two values of  $\ell$ .

Complementing Figure 3.5, Figure 3.6 shows the relative increase,  $R$  say, in regret for six different methods, with  $R$  defined as  $R = (a - b)/b$ , where  $a$  equals the regret when the empirically chosen version of  $h_1$  is used, and  $b$  is the regret for the optimal choice of  $h_1$ . This comparison shows that the rule based on  $\widehat{\text{err}}_{\mathcal{A}_1}$  performs better than the other five approaches; that cross-validation and the 0.632+ bootstrap perform worst; and that bagged cross-validation, bagged two-fold cross-validation and Efron's (1983) approach are between those two groups.

We carried out the same experiment for the case where the normal mixture is replaced by a mixture of two lognormal distributions, or a mixture of two Cauchy distributions. In each setting the components in the mixture were taken to have the same location and scale parameters as in the normal case. The results are presented in Figure 3.7, and closely reflect those in Figure 3.6. In particular, apart from the Cauchy mixture case with  $n = 50$ , cross-validation and 0.632+ bootstrap methods again give the highest regret ratios. In the Cauchy case, although generally not for lognormal data, Efron's bootstrap methods edges out the bootstrap approach suggested in Section 2.

We also explored properties of cross-validation in problems where at least one of the densities  $f$  and  $g$  becomes increasingly complex as sample size increases.

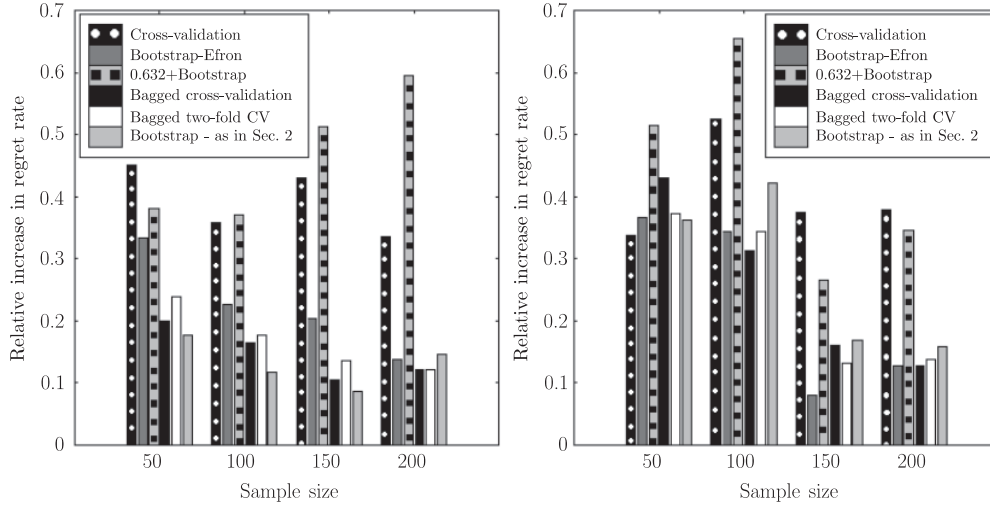


Figure 3.7. Relative increase in regret for different classifiers, in the case of (a) lognormal mixtures or (b) Cauchy mixtures.




Figure 3.8. Relative increase in regret for different choices of  $n$  and  $k$ , when the problem becomes increasingly complex with sample size.

This setting favours cross-validation. For example, taking  $f$  to be the uniform density on  $[0, 1]$  and  $g(x) = 1 + \cos(2k\pi x)$ , the value of  $R$ , in the case of cross-validation, decreases as  $k$  increases across a broad range; see Figure 3.8. The ratio also decreases if  $k$  and  $n$  increase together, in particular if  $k = \log n$ .

This reflects the fact that cross-validation is essentially a global procedure;

it performs well at estimating tuning parameters determined by global issues, but does poorly at estimating those parameters when they have only a local influence. If two densities cross at only a small number of points then, since the performance of a classifier is determined by properties of the densities close to those points, optimising the classifier is a distinctly local problem. Therefore, cross-validation performs poorly. However, as the number of crossing points increases, the problem becomes more global in nature, and cross-validation becomes more competitive.

## References

- Ancukiewicz, M. (1998). An unsupervised and nonparametric classification procedure based on mixtures with known weights. *J. Classification* **15**, 129-141.
- Bowman, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika* **71**, 353-360.
- Breiman, L. (1996). Bagging predictors. *Mach. Learn.* **24**, 123-140.
- Chanda, K. C. and Ruymgaart, F. H. (1989). Asymptotic estimate of probability of misclassification for discriminant rules based on density estimates. *Statist. Probab. Lett.* **8**, 81-88.
- Devroye, L., Györfi, L. and Lugosi, G. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer, New York.
- Efron, B. (1983). Estimating the error rate of a prediction rule: improvement on cross-validation. *J. Amer. Statist. Assoc.* **78**, 316-331.
- Efron, B. and Tibshirani, R. (1997). Improvements on cross-validation: the .632+ bootstrap method. *J. Amer. Statist. Assoc.* **92**, 548-560.
- Faraway, J. J. and Jhun, M. (1990). Bootstrap choice of bandwidth for density estimation. *J. Amer. Statist. Assoc.* **85**, 1119-1122.
- Fix, E. and Hodges, J. (1951). Discriminatory analysis. Nonparametric discrimination: consistency properties. Technical Report No. 4, Project No. 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX.
- Ghosh, A. K. and Chaudhuri, P. (2004). Optimal smoothing in kernel discriminant analysis. *Statist. Sinica* **14**, 457-483.
- Ghosh, A. K. and Hall, P. (2006). On error-rate estimation in nonparametric classification. Manuscript.
- Hall, P. (1983). Large sample optimality of least squares cross-validation in density estimation. *Ann. Statist.* **11**, 1156-1174.
- Hall, P. and Kang, K.-H. (2005). Bandwidth choice for nonparametric classification. *Ann. Statist.* **33**, 284-306.
- Krzyżak, A. (1991). On exponential bounds on the Bayes risk of the nonparametric classification rules. In *Nonparametric Functional Estimation and Related Topics* (Spetses, 1990), *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.* **335**, 347-360.
- Lapko, A. V. (1993). *Nonparametric Classification Methods and their Application*. (In Russian.) VO Nauka, Novosibirsk.
- Lin, C. T. (2001). Nonparametric classification on two univariate distributions. *Commun. Statist. Theory Meth.* **30**, 319-330.

- Lugosi, G. and Pawlak, M. (1994). On the posterior-probability estimate of the error rate of nonparametric classification rules. *IEEE Trans. Inform. Theory* **40**, 475-481.
- Lugosi, G. and Nobel, A. (1996). Consistency of data-driven histogram methods for density estimation and classification. *Ann. Statist.* **24**, 687-706.
- Mammen, E. and Tsybakov, A. B. (1999). Smooth discrimination analysis. *Ann. Statist.* **27**, 1808-1829.
- Pawlak, M. (1993). Kernel classification rules from missing data. *IEEE Trans. Inform. Theory* **39**, 979-988.
- Steele, B. M. and Patterson, D. A. (2000). Ideal bootstrap estimation of expected prediction error for  $k$ -nearest neighbor classifiers: Applications for classification and error assessment. *Statist. Comput.* **10**, 349-355.
- Stone, C. J. (1984). An asymptotically optimal window selection rule for kernel density estimates. *Ann. Statist.*, 1285-1297.
- Yang, Y. H. (1999a). Minimax nonparametric classification — Part I: Rates of convergence. *IEEE Trans. Inform. Theory* **45**, 2271-2284.
- Yang, Y. H. (1999b). Minimax nonparametric classification — Part II: Model selection for adaptation. *IEEE Trans. Inform. Theory* **45**, 2285-2292.

Theoretical Statistics and Mathematics Unit, Indian Statistical Institute, Kolkata 700108, India.

E-mail: anilghosh@rediffmail.com

Department of Mathematics and Statistics, University of Melbourne, Melbourne, VIC 3010, Australia.

E-mail: halpstat@ms.unimelb.edu.au

(Received June 2006; accepted January 2007)