Cross Validation for Comparing Multiple Density Estimation Procedures

Heng Lian

Division of Applied Mathematics
Brown University
182 George St
Providence, RI 02912

Abstract

We demonstrate the consistency of cross validation for comparing multiple density estimators using simple inequalities on the likelihood ratio. In nonparametric problems, the splitting of data does not require the domination of test data over the training/estimation data, contrary to Shao (1993). The result is complementary to that of Yang (2005) and Yang (2006).

Key words: Density estimation, Cross validation, Kullback-Leibler divergence, Hellinger distance

1 Introduction

Cross validation (CV) is a common procedure used for smoothing parameter estimation or model selection. When multiple procedures are being compared, it is customary to choose the one that obtains the smallest "loss", which is defined specifically to the problem at hand. If both the procedure used to obtain the estimate and the computation of the loss are based on the same set of data, it is a well-known effect that the estimated loss is biased due to the double use of the same observations. In order to obtain an unbiased estimate, one approach is to use a penalty term that takes into account the complexity of the model. This approach includes AIC, BIC, C_p, etc. A simpler approach, which is closely related, when we have the luxury of enough observations, is to split the data in such way that one part is used to obtain the estimate, and the other separate hold-out data is used to evaluate the loss. The main advantage

---

of this approach is that it can be easily applied in various situations (without theoretical derivation) to select one out of many competing procedures. There are a few different variations to this approach, including leave-one-out CV, k-fold CV, Generalized CV, etc. A well-studied case is the linear regression problem, where Shao (1993) showed the surprising result that leave-one-out CV will select the models with extra redundant variables with nonvanishing probability (This theory assumed that the number of covariates is fixed. It is a different story when the number of covariates grows with n). In order to restore consistency, one should split the data such that the size of the evaluation part of the data is dominating. All of the above techniques are summarized and compared in Shao (1997) in the context of linear regression with different kinds of asymptotics.

Yang (2005) studied the problem of cross-validation in the context of non-parametric regression comparing a finite number of estimators. It is shown in that paper that under the L_2 loss, as long as one of the competing procedures converges at a nonparametric rate, the dominance of evaluation data is not necessary for consistency. Instead, the two parts of the data can be of the same order, which is surprising considering the corresponding result for linear regression.The proof of Yang's result is based on an application of Bernstein's inequality. Similarly, it is shown in Yang (2006) that cross validation is consistent in classification problems where the consistency also depends on the rate of disagreement between the two classifiers.

In this paper, we consider the problem of density estimation when the observations are generated i.i.d. from the true distribution P_0. There exists a large literature on density estimators, earlier results focus on linear estimators including kernel density estimator, later developments include wavelets thresholding and adaptive width kernel that achieve minimax rate of convergence in a large class of Besov spaces where no linear estimate can attain the optimal convergence rate. Faced with such large choices of estimator with different theoretical properties, it is important to select one that has optimal performance for the current problem. Cross validation can be directly applied by splitting the observations into two groups. Different estimators, such as kernel estimates and wavelets, can be obtained based on the first part of data, then the likelihood of the second group of data can be evaluated and compared. Finally, the estimator that obtains the largest likelihood is chosen as the winner. A natural question is whether this process will return the optimal procedure. In particular, what condition on the splitting ratio should be satisfied in order to ensure the consistency property? The main conclusion in this paper is similar to that of Yang (2005), that is, we do not necessarily need to assume the size dominance of the evaluation data as in linear regression problem.

2 Consistency result

Consider the situation where we have n observations $X_1, X_2, \ldots, X_n$ generated i.i.d. from an underlying distribution P_0 with corresponding density p_0. There exists many well-known density estimators. For parametric procedures, such as mixture modeling using parametric families of densities, the convergence rate is usually $1/\sqrt{n}$. On the other hand, for nonparametric procedures, the rate of convergence is slower, depending on the smoothness property of the true density. Density estimation is closely related to regression problem, as demonstrated in a series of well-known papers (see, e.g., Brown and Low (1996), Nussbaum (1996)).

With many possible choices for density estimation procedures, both parametric and nonparametric, one needs to find the best estimator among them for the current data. Parametric procedures have a faster rate of convergence when the model is correct, but suffer from a nonvanishing bias when the true distribution lies outside of the parametric family. Nonparametric procedures are more flexible but lose in efficiency when the underlying density is of a known parametric form. In practice, we need to know which procedure is best without knowledge of the true distribution.

We start by splitting the data into two parts: the estimation data $X^1 = \{X_1, \ldots, X_{n_1}\}$ and the evaluation data $X^2 = \{X_{n_1+1}, \ldots, X_n\}$, and let $n_2 = n - n_1$. We assume we have many estimation procedures $\{\hat{P}_i\}_{i=1}^{m_n}$ (note the number of potential choices can grow with n), which will produce density estimates $\{\hat{p}_i^{(n_1)}(x; X_1, \ldots, X_{n_1})\}$, we will omit the dependence of $\hat{p}_i^{(n_1)}$ on the training data X^1 in the following. To choose the best procedure among those m_n, the test data X^2 is used to evaluate the likelihood: $\hat{p}_i^{(n_1)}(X^2) = \prod_{k=n_1+1}^n \hat{p}_i^{(n_1)}(X_k)$. If $\hat{p}_j^{(n_1)}(X^2) = \max_i \hat{p}_i^{(n_1)}(X^2)$, then the procedure $\hat{P}_j^{(n_1)}$ is selected as the final estimator. The desired property is that this cross-validation procedure will select the best one with high probability. We will use a loss function $d(p_0, p)$ to measure the closeness of p to the true density p_0. In this paper, we will adopt the commonly used Hellinger distance as the loss function: $d_H(p_0, p) = (\int (\sqrt{p_0} - \sqrt{p})^2)^{1/2}$. Another commonly used measure of loss in the context of density estimation is the Kullback-Leibler divergence (which is not a true distance) $d_K(p_0, p) = \int p_0 \log \frac{p_0}{p}$. It is always true that $d_H^2 \leq d_K$. Under some mild assumptions, these two measures of loss are almost equivalent. The simplest case under which this is true is when the class of densities considered are uniformly bounded away from zero and infinity, so that the ratio $\frac{p_0}{p}$ is uniformly bounded, then $d_K(p_0, p) = O(d_H^2(p_0, p))$ (see, e.g., Lemma 8.2 in Ghosal et al. (2000)). The more complicated techniques similar to those used in section 3 of that paper can also be used when the estimate is constrained to be within a finite approximation set (this will result in an extra logarithmic factor). The paper of Yang and Barron (1999) contains more infor-

mation regarding the relationships between d_H^2, d_K and the L_p loss function, and established some equivalence result between them under some conditions. In the rest of the paper, we will assume that $d_K \leq M d_H^2$ for some constant M.

Theorem 1 $P_0(\hat{p}_1^{(n_1)}(X^2) > \hat{p}_i^{(n_1)}(X^2), \forall i > 1 | X^1) \to 1$ *if the following conditions hold:*

(1) $n_1 \to \infty, n_2 \to \infty$

(2) $n_2 \min_{i>1} v_{n_1,i}^2 \to \infty$, *and* $\log m_n = o(n_2 \min_{i>1} v_{n_1,i}^2)$

(3) There exists $c < 1$ *and* $t_n > 0$ *s.t.* $nt_n \to \infty$ *and* $\frac{M v_{n_1,1}^2 + s_{n_1} t_{n_2}}{c v_{n_1,i}^2} < 1, \forall i > 1$

where $v_{n,i} = d_H(p_0, \hat{p}_i^{(n)})$, $s_n = V(p_0, \hat{p}_1^{(n)})$, *and* $V(p_0, p) = \int p_0 (\log \frac{p_0}{p})^2$

Remark 2 *In the statement of the theorem, the probability is conditioned on* X^1, *and* $v_{n_1,i}$ *and* s_{n_1} *in conditions (2) and (3) are random variables, so the theorem should be interpreted as* $P_0(\hat{p}_1^{(n_1)}(X^2) > \hat{p}_i^{(n_1)}(X^2), \forall i > 1 | X^1) \to 1$ *on the set that the conditions (1)-(3) hold.*

Proof. The proof is based on simple likelihood ratio inequalities in Wong and Shen (1995). All the probabilities below are implicitly conditioned on the training data X^1. From Lemma 1 in Wong and Shen (1995), we have, for $i \neq 1, b > 0$,

$$P_0(\frac{\hat{p}_i^{(n_1)}(X^2)}{p_0(X^2)} \geq exp(-n_2 b)) \leq exp(\frac{n_2 b}{2} - \frac{n_2 v_{n_1,i}^2}{2})$$

Choosing $b = c v_{n_1,i}^2$ (c as in the above assumption (3)), we get

$$P_0(\frac{\hat{p}_i^{(n_1)}(X^2)}{p_0(X^2)} \geq exp(-n_2 c v_{n_1,i}^2)) \leq exp(-\frac{n_2 v_{n_1,i}^2}{2}(1 - c))$$

Denote by W_{n_2} the event $\{\frac{d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) - d_K(p_0, \hat{p}_1^{(n_1)})}{V(p_0, \hat{p}_1^{(n_1)})} \geq t_{n_2}\}$, where $d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)})$ is the empirical version of $d_K(p_0, \hat{p}_1^{(n_1)})$ on evaluation data X^2:

$$d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) = \frac{1}{n_2} \sum_{i=n_1+1}^{n} \log \frac{p_0(X_i)}{\hat{p}_1^{(n_1)}(X_i)}$$

. By Chebyshev's inequality, $P_0(W_{n_2}) \leq \frac{1}{n_2 t_{n_2}}$. Denoting $d = M v_{n_1,1}^2 + s_{n_1} t_{n_2}$,

4

$$P_0(\frac{\hat{p}_1^{(n_1)}(X^2)}{p_0(X^2)} \le exp(-n_2 d))$$

$$= P_0(exp\{-n_2 V(p_0, \hat{p}_1^{(n_1)})\frac{d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) - d_K(p_0, \hat{p}_1^{(n_1)})}{V(p_0, \hat{p}_1^{(n_1)})}\}exp\{-n_2 d_K(p_0, \hat{p}_1^{(n_1)})\}$$

$$\le exp\{-n_2 d\})$$

$$\le P_0\left(exp\{-n_2 V(p_0, \hat{p}_1^{(n_1)})\frac{d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) - d_K(p_0, \hat{p}_1^{(n_1)})}{V(p_0, \hat{p}_1^{(n_1)})}\} \le exp\{-n_2(d - Mv_{n_1,1}^2)\}\right)$$

$$\le P_0(W_{n_2}) \to 0$$

The last inequality above holds since on the set

$$W_{n_2}^c = \{\frac{d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) - d_K(p_0, \hat{p}_1^{(n_1)})}{V(p_0, \hat{p}_1^{(n_1)})} < t_{n_2}\},$$

we have

$$exp\{-n_2 V(p_0, \hat{p}_1^{(n_1)})\frac{d_K^{(n_2)}(p_0, \hat{p}_1^{(n_1)}) - d_K(p_0, \hat{p}_1^{(n_1)})}{V(p_0, \hat{p}_1^{(n_1)})}\} > exp\{-n_2 s_{n_1} t_{n_2}\}$$

$$= exp\{-n_2(d - Mv_{n_1,1}^2)\}$$

Note that when $d = Mv_{n_1,1}^2 + s_{n_1} t_{n_2} < cv_{n_2,i}^2$ (assumption (3)), $\frac{\hat{p}_1^{(n_1)}(X^2)}{p_0(X^2)} > e^{-n_2 d}$ and $\frac{\hat{p}_i^{(n_1)}(X^2)}{p_0(X^2)} < e^{-n_2 cv_{n_2,i}^2}$ implies $\hat{p}_1^{(n_1)}(X^2) > \hat{p}_i^{(n_1)}(X^2)$. So we can bound the probability that the cross validation procedure chooses an estimator other than $\hat{P}_1$:

$$P_0(\hat{p}_1^{(n_1)}(X^2) < \hat{p}_i^{(n_1)}(X^2) \text{ for some } i)$$

$$\le P_0(\frac{\hat{p}_1^{(n_1)}(X^2)}{p_0(X^2)} < e^{-n_2 d}) + \sum_{i=2}^{m_n} P_0(\frac{\hat{p}_i^{(n_1)}(X^2)}{p_0(X^2)} > e^{-n_2 cv_{n_2,i}^2})$$

$$\le P_0(W_{n_2}) + m_n exp(-\frac{n_2(1-c)}{2} \min_i v_{n_1,i}^2)$$

The above expression converges to zero under the assumed condition (2). □

Remark 3 *The above theorem cannot be directly applied since $v_{n_1,i}$ are random variables depending on X^1. We will specialize the result to the two procedures case below in Corollary 1.*

Remark 4 *Under some mild conditions, we will have $s_{n_1} = V(p_0, \hat{p}_1^{(n_1)}) = O(d_H^2(p_0, \hat{p}_1^{(n_1)}))$, see, e.g., Theorem 5 in Wong and Shen (1995).*

Remark 5 *As in Yang (2005), we can consider the case where we have multiple different splittings of the original data. Cross validation as stated above can be applied to each splitting separately, and then use majority vote to choose the final procedure.*

Now we give a definition comparing two procedures by their rate of convergence.

Definition 6 *Procedure $\hat{P}_1$ is asymptotically better than $\hat{P}_2$ under the loss function d_H if for some sequence $\epsilon_n \to 0$*

$$\lim_{n \to \infty} P_0(d_H(p_0, \hat{p}_1^{(n)}) < \epsilon_n d_H(p_0, \hat{p}_2^{(n)})) \to 1$$

Under this definition we can state the following corollary, the definition of exact rate of convergence is similar to Definition 3 in Yang (2005).

Corollary 7 *Considering two procedures for density estimation where one is asymptotically better than the other. Suppose the exact rate of convergence of $d_H(p_0, \hat{p}_1^{(n)})$ and $d_H(p_0, \hat{p}_2^{(n)})$ are p_n and q_n respectively. Assume that $V(p_0, \hat{p}_i^{(n_1)}) = O(d_K(p_0, \hat{p}_i^{(n_1)})), i = 1, 2$. If $n_1 \to \infty, n_2 \to \infty, \sqrt{n_2} \max(p_{n_1}, q_{n_1}) \to \infty$, then the cross validation is consistent in the sense of choosing the asymptotically better procedure with probability tending to 1.*

Remark 8 *If one of the procedures has nonparametric rate $n^{-\alpha}$ with $\alpha < 1/2$, then $\frac{n_1}{n_2} = O(1)$ will suffice for $\sqrt{n_2} \max(p_{n_1}, q_{n_1}) \to \infty$*

3   Conclusions

We give a simple proof of the consistency of cross validation in the context of density estimation. Although it is shown in Shao (1993) that leave-one-out cross validation is inconsistent for linear regression problem, it is unclear to us whether this is the case for nonparametric problems. Another interesting problem is that when multiple splittings are available, we can either use majority voting as in Yang (2005) or choose the procedure with the largest product of individual likelihood for each splitting. The comparison of these two approaches is similar to the tradeoff between model selection and model averaging.

References

Brown,L.D., Low,M.G., 1996. Asymptotic equivalence of nonparametric regression and white noise. Ann. Statist. 24, 2384-2398.

Ghosal,S., Ghosh,J.K., van der Vaart,A.W., 2000. Convergence rates of posterior distributions, Ann. Statist. 28, 500-531.

Nussbaum,M., 1996. Asymptotic equivalence of density estimation and Gaussian white noise, Ann. Statist. 24, 2399-2430.

Shao,J., 1993. Linear model selection by cross-validation, J. Amer. Statist. Assoc. 88, 486-494.

Shao,J., 1997. An asymptotic theory for linear model selection, Statist. Sinica 7, 221-264.

Wong,W.H., Shen,X., 1995. Probability inequalities for likelihood ratios and convergence rates of sieve MLEs, Ann. Statist. 23, 339-362.

Yang,Y., 2005. Consistency of Cross Validation for Comparing Regression Procedures, submitted to Ann. Statist.

Yang,Y., 2006. Comparing Learning Methods for Classification, Statist. Sinica 16, 635-657.

Yang,Y., Barron,A., 1999. Information-theoretic determination of minimax rates of convergence, Ann. Statist. 27, 1564-1599.