

CV-NP BAYESIANISM BY MCMC

Cross Validated Non Parametric Bayesianism by Markov Chain Monte Carlo

CARLOS C. RODRIGUEZ
Department of Mathematics and Statistics
University at Albany, SUNY
Albany NY 12222, USA[†]

Abstract. Completely automatic and adaptive non-parametric inference is a pie in the sky. The frequentist approach, best exemplified by the kernel estimators, has excellent asymptotic characteristics but it is very sensitive to the choice of smoothness parameters. On the other hand the Bayesian approach, best exemplified by the mixture of gaussians models, is optimal given the observed data but it is very sensitive to the choice of prior. In 1984 the author proposed to use the Cross-Validated gaussian kernel as the likelihood for the smoothness scale parameter h , and obtained a closed formula for the posterior mean of h based on Jeffreys's rule as the prior. The practical operational characteristics of this bayes' rule for the smoothness parameter remained unknown for all these years due to the combinatorial complexity of the formula. It is shown in this paper that a version of the metropolis algorithm can be used to approximate the value of h producing *remarkably* good completely automatic and adaptive kernel estimators. A close study of the form of the cross validated likelihood suggests a modification and a new approach to Bayesian Non-parametrics in general.

Key words: Cross validation, Density estimation, Bayes method, Kernel estimator

1. Introduction

A basic problem of statistical inference is to estimate the probability distribution that generated the observed data. In order to allow the data to *speak by themselves*, it is desirable to solve this problem with a minimum of a priori assumptions on the class of possible solutions. For this reason, the last thirty years have been burgeoning with interest in nonparametric estimation. The main positive results are almost exclusively from non-Bayesian quarters, but due to recent advances in Monte Carlo methods and computer hardware, Bayesian nonparametric estimators are now becoming more competitive.

[†]Email: carlos@math.albany.edu

This paper aims to extend the idea in [1] to a new general Bayesian approach to nonparametric density estimation. As an illustration one of the techniques is applied to compute the estimator in [1]. In particular it is shown how to approximate the bayes estimator (for Jeffreys' prior and quadratic loss) for the bandwidth parameter of a kernel using a version of the metropolis algorithm. The computer experiments with simulated data clearly show that the bayes estimator with Jeffreys' prior outperforms the standard estimator produced by plain Cross-validation.

2. Density Estimation by Summing Over Paths

Any non-Bayesian nonparametric density estimator can be turned into a Bayesian one by the following three steps:

1. Transform the estimator into a likelihood for the smoothness parameters via the sum over paths technique introduced below.
2. Put a reference prior on the smoothness parameters.
3. Use the predictive distribution as the Bayesian estimator obtainable by Markov Chain Monte Carlo.

The summing over paths method, springs from the simple idea of interchanging sums and products in the expression for the cross-validated kernel likelihood (see [1]). Without further reference to its humble origins, let us postulate the following sequence of functions,

$$\Phi_n = \Phi(x_0, x_1, x_2, \dots, x_n | h) \propto \sum_{\text{all paths}} \prod_{j=0}^n K_h(x_j - x_{i_j}) \quad (1)$$

where $h > 0$, $K_h(x) = \frac{1}{h}K(x/h)$ with K a density symmetric about 0. We call a vector of indices, (i_0, \dots, i_n) with the property that $i_j \in \{0, \dots, n\} \setminus \{j\}$, a path (more specifically a general unrestricted path, see below). The sum above runs over all the n^{n+1} possible paths. The functions, Φ_n are defined up to a proportionality constant independent of the x_j 's.

Notice that by flipping the sum and the product we get

$$\frac{1}{n^{n+1}} \sum_{\text{all paths}} \prod_{j=0}^n K_h(x_j - x_{i_j}) = f_{-0,n}(x_0) f_{-1,n}(x_1) \cdots f_{-n,n}(x_n) \quad (2)$$

where,

$$f_{-j,n}(x_j) = \frac{1}{n} \sum_{\substack{i_j=0 \\ i_j \neq j}}^n K_h(x_j - x_{i_j}). \quad (3)$$

Thus, $f_{-j,n}(x_j)$ is nothing but the kernel density estimator of $f(x_j)$ based on all the data except the j th observation x_j . Under mild regularity conditions the kernel estimator is known to converge (in every conceivable way) provided that $h = h_n$ is taken as a function of n such that, $h_n \rightarrow 0$ and $nh_n \rightarrow \infty$ as $n \rightarrow \infty$.

The Φ_n 's can be used as a *universal* method for attaching a class of exchangeable one parameter models to any set of observations. The positive scalar parameter h is the only free parameter, and different models are obtained by changing the kernel function K .

These empirical parametric models are invariant under relabeling of the observations (i.e. they are exchangeable) but they do not model the observations as independent variables. Rather, these models introduce a pattern of correlations for which there is a priori no justification. This suggest that there might be improvements in performance if the sum is restricted to special subsets of the set of all n^{n+1} paths. Three of these modifications are mentioned in the following section.

Notice also that the ability of the Φ_n to adapt comes at the expense of regularity. These models are always non-regular. If the kernel has unbounded support then Φ_n integrates to infinity but the conditional distribution of x_0 given x_1, \dots, x_n and h is proper. When the kernel has compact support the Φ_n are proper but still non-regular since their support now depends on h .

The above recipe would have been a capricious combination of symbols if it not were for the fact that, under mild regularity conditions, these models *adapt* to the form of the true likelihood as n increases.

As a function of $x_0 = x$, the Φ_n have the following asymptotic property,

Theorem 1 *If x_1, x_2, \dots, x_n are iid observations from an unknown pdf f which is continuous a.s. and $h = h_n$ is taken as a function of n such that, $h_n \rightarrow 0$ and $nh_n \rightarrow \infty$ as $n \rightarrow \infty$, then,*

$$\frac{\Phi_n}{\int \Phi_n dx_0} = f_{-0,n}(x) + o(1) = f(x) + o(1). \quad (4)$$

where the little o is taken in probability as $n \rightarrow \infty$

Proof (sketch only)

Just flip the sum and the product to get again,

$$\frac{1}{n^{n+1}} \sum_{\text{all paths}} \prod_{j=0}^n K_h(x_j - x_{i_j}) = f_{-0,n}(x) f_{-1,n}(x_1) \cdots f_{-n,n}(x_n) \quad (5)$$

Under the simple regularity conditions of the theorem, the kernel estimator is known to converge in probability as $n \rightarrow \infty$. However, even though x_0 appears in all of the $n + 1$ factors, and their number goes to infinity, still all the factors are converging to the value of the true density at the given point. Therefore the theorem follows. \square

It is worth noticing that the above theorem is only one of a large number of results that are readily available from the thirty years of literature on density estimation. In fact under appropriate regularity conditions the convergence can be strengthened to be pointwise a.s., uniformly a.s., or globally a.s. in L_1 , or L_2 .

3. Paths, Graphs and Loops

Each of the n^{n+1} paths (i_0, \dots, i_n) can be represented by a graph with nodes x_0, \dots, x_n and edges from x_j to x_k if and only if $i_j = k$. Here are some graphs

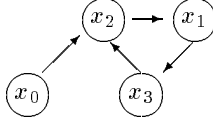


Figure 1. The graph of $(2, 3, 1, 2)$

for paths with $n = 3$. For example, the path $(2, 3, 1, 2)$ is given a probability proportional to

$$K_h(x_0 - x_2)K_h(x_1 - x_3)K_h(x_2 - x_1)K_h(x_3 - x_2) \quad (6)$$

and represented by the graph in figure [1]. Let's call it a 1-3-loop. The path $(1, 2, 3, 0)$ is the single ordered loop of size four (a 4-loop), $(3, 0, 1, 2)$ is the same loop backwards (also a 4-loop), $(2, 3, 0, 1)$ are two disconnected loops (a 2-2-loop) and $(1, 0, 0, 0)$ is connected and contains a loop of size two with x_0 and x_1 in it (a 1-1-2-loop). Draw the pictures!

The classification of paths in terms of number and size of loops appears naturally when trying to understand how Φ_n distributes probability mass among the different paths. To be able to provide simple explicit formulas let us take K in the definition of Φ_n to be the standard gaussian density, i.e. from now on we take,

$$K(x) = \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-x^2}{2}\right). \quad (7)$$

The gaussian kernel has unbounded support and that makes the total integral of each path to diverge. Thus, the partition function

$$Z = \int \Phi_n dx_0 \dots dx_n \quad (8)$$

$$= \sum_{\text{all paths}} \int \prod_{j=0}^n K_h(x_j - x_{i_j}) dx_0 \dots dx_n \quad (9)$$

is the sum of infinities and it also diverges. Recall that this anomaly is the price we need to pay for using a model with a finite number of free parameters (only one in this case) and hoping to still adapt to the form of the true likelihood as $n \rightarrow \infty$. Even though the value of Z is in fact ∞ we can still write (formally) a useful decomposition that will help explain how the Φ_n 's adapt and how to modify the set of paths to improve the convergence. We first need the following simple

property of gaussians,

$$\int K_a(x-y)K_b(y-z)dy = K_{\sqrt{a^2+b^2}}(x-z). \quad (10)$$

This can be shown by straight forward integration after completing the square. Now notice that whatever the value of the integrals appearing in equation (9) that value only depends on the type of loop that is being integrated. For this reason we omit the integrand and simply denote the value of the integral with the integral sign and the type of loop. With this notation we have,

Theorem 2

$$\int_{m_1-m_2-\dots-m_k\text{-loop}} = \left(\int_{m_1\text{-loop}} \right) \left(\int_{m_2\text{-loop}} \right) \dots \left(\int_{m_k\text{-loop}} \right) \quad (11)$$

More over, $\int_{1\text{-loop}} = 1$ and for $m > 1$,

$$\int_{m\text{-loop}} = \frac{1}{\sqrt{2\pi}} \frac{h^{-1}}{\sqrt{m}} L \quad (12)$$

where we write formally $L = \int dx$.

Proof

Equation (11) follows from Fubini's theorem. To get (12) use Fubini's theorem and apply (10) each time to obtain,

$$\begin{aligned} \int_{m\text{-loop}} &= \int K_h(x_0-x_1)K_h(x_1-x_2)K_h(x_2-x_3)\dots K_h(x_{m-1}-x_0)dx_0dx_1\dots dx_{m-1} \\ &= \int K_{\sqrt{2}h}(x_0-x_2)\dots K_h(x_{m-1}-x_0)dx_0dx_2\dots dx_{m-1} \\ &\dots \\ &= \int K_{\sqrt{m-1}h}(x_0-x_{m-1})K_h(x_{m-1}-x_0)dx_0dx_{m-1} \\ &= \int K_{\sqrt{m}h}(0)dx_0 = \frac{1}{\sqrt{2\pi}} \frac{h^{-1}}{\sqrt{m}} L \end{aligned}$$

□

Hence, by splitting the sum over all paths into,

$$\sum_{\text{all paths}} = \sum_{2-2\dots-2\text{-loops}} + \sum_{1-3-2\dots-2\text{-loops}} + \dots + \sum_{(n+1)\text{-loops}}$$

and applying the previous theorem we obtain,

$$Z = N_{2-2\dots-2} \left(\frac{1}{\sqrt{2\pi}} \frac{h^{-1}}{\sqrt{2}} L \right)^{(n+1)/2} + \dots + N_{n+1} \left(\frac{1}{\sqrt{2\pi}} \frac{h^{-1}}{\sqrt{n+1}} L \right)^1 \quad (13)$$

where for simplicity we have assumed that n is odd and we denote by $N_{m_1-\dots-m_k}$ the total number of $m_1-\dots-m_k$ -loops. Using simple combinatorial arguments

it is possible to write explicit formulas for the number of loops of each kind. The important conclusion from the decomposition (13) is that even though the Φ_n appear to be adding equally over all paths, in reality they end up allocating almost all the probability mass on paths with maximally disconnected graphs. This is not surprising. This is the reason why there is consistency when assuming iid observations. There is a built in bias towards independence. The bias can be imposed explicitly on the Φ_n by restricting the paths to be considered in the sum. Here are three examples:

loops:

Only paths (i_0, \dots, i_n) that form a permutation of the integers $\{0, 1, \dots, n\}$ are considered.

2-2-...-2-loops:

Only maximally disconnected paths are considered.

QM:

Paths as above but use $|\Phi_n|^2$ instead of Φ_n as the joint likelihoods.

Preliminary simulation experiments seem to indicate that only maximally disconnected paths are not enough and that all the loops are too many. The QM method has all the maximally disconnected paths but not all the loops (e.g. with $n = 5$ the 3-3-loops can not be reached by squaring the sum of 2-2-2-loops) so it looks like the most promising among the three. What is more interesting about the QM method is the possibility of using kernels that can go negative or even complex valued. More research is needed since very little is known about the performance of these estimators.

4. Estimation by MCMC

We show in this section how to approximate the predictive distribution and the bayes rule for the smoothness parameter by using Markov Chain Monte Carlo techniques.

4.1. POSTERIOR MEAN OF THE BANDWIDTH

Apply bayes' theorem to (1) to obtain the posterior distribution of h ,

$$\pi(h|x, x_1, x_2, \dots, x_n) = \frac{\Phi(x, x_1, x_2, \dots, x_n|h)\pi(h)}{\int_0^\infty \Phi(x, x_1, x_2, \dots, x_n|\tau)\pi(\tau)d\tau} \quad (14)$$

where π is a function of h . It is worth noticing that π is **not** the prior on h . It is only the part of the prior on h that we can manipulate. Recall that Φ_n integrates to the function of h given by (13) so effectively the prior that is producing (14) is,

$$\Pi(h) \propto \frac{\pi(h)}{h^{(n+1)/2}} \quad (15)$$

The posterior mean is then given by,

$$E(h|x, x_1, x_2, \dots, x_n) = \hat{h}_x = \frac{\int_0^\infty h\Phi(x, x_1, x_2, \dots, x_n|h)\pi(h)dh}{\int_0^\infty \Phi(x, x_1, x_2, \dots, x_n|h)\pi(h)dh} \quad (16)$$

Equation (16) provides a different estimator for each value of x . To obtain a single global estimate for h just erase the x 's from (16) and change n to $n - 1$ in the formulas below. When K is the univariate gaussian kernel and $\pi(h) = h^{-\delta}$ equation (16) simplifies to:

$$\hat{h}_x = C_{n,\delta} \frac{\sum_{\text{all paths}} \alpha(\underline{i}) s(\underline{i})}{\sum_{\text{all paths}} \alpha(\underline{i})} \quad (17)$$

where,

$$C_{n,\delta} = \frac{1}{\sqrt{2}} \frac{\Gamma\left(\frac{n+\delta-1}{2}\right)}{\Gamma\left(\frac{n+\delta}{2}\right)} \quad (18)$$

$\underline{i} = (i_0, \dots, i_n)$ is a path, $\alpha = s^{-(n+\delta)}$ and

$$s^2(\underline{i}) = \sum_{j=0}^n (x_j - x_{i_j})^2. \quad (19)$$

Equation (17) follows from two applications of the formula,

$$\int_0^\infty h^{-(\beta+1)} \exp\left\{-\frac{s^2}{2h^2}\right\} dh = \frac{1}{2} 2^{\beta/2} \Gamma(\beta/2) s^{-\beta} \quad (20)$$

4.1.1. Bandwidth by Metropolis

To approximate equation (17) we use the fact that the ratio of the two sums is the expected value of a random variable that takes the value $s(\underline{i})$ on the path \underline{i} which is generated with probability proportional to $\alpha(\underline{i})$. The following version of the Metropolis algorithm produces a sequence of averages that converge to the expected value,

Algorithm

0) Start from somewhere

$$\begin{aligned} \underline{i} &\leftarrow (1, 2, \dots, n, 0) \\ s2 &\leftarrow \sum_{j=0}^n (x_j - x_{i_j})^2 \\ \alpha &\leftarrow (s2)^{-(n+\delta)/2} \\ N &\leftarrow 0, \text{ sum} \leftarrow 0, \text{ ave} \leftarrow 0 \end{aligned}$$

1) Sweep along the components of \underline{i}
for k from 0 to n do

$$\left\{ \begin{aligned} i'_k &\leftarrow \text{Uniform on } \{0, \dots, n\} \setminus \{k, i_k\} \\ \Delta_k &\leftarrow (x_{i'_k} - x_{i_k})(x_{i'_k} + x_{i_k} - 2x_k) \\ s2' &\leftarrow s2 + \Delta_k \\ \alpha' &\leftarrow (s2')^{-(n+\delta)/2} \\ R &\leftarrow \alpha'/\alpha \end{aligned} \right.$$

if $R > 1$ or $\text{Unif}[0,1] < R$ then $\{i_k \leftarrow i'_k, s2 \leftarrow s2', \alpha \leftarrow \alpha'\}$
 $\}$

2) Update the estimate for the average,
 $\text{sum} \leftarrow \text{sum} + \sqrt{s2}$
 $N \leftarrow N + 1$
 $\text{ave} \leftarrow \text{sum}/N$
 goto 1)

4.2. THE PREDICTIVE DISTRIBUTION BY GIBBS

To sample from the predictive distribution, $f(x|x_1, x_2, \dots, x_n)$ we use Gibbs to sample from the joint distribution, $f(x, h|x_1, x_2, \dots, x_n)$. Hence, we only need to know how to sample from the two conditionals, a) $f(x|h, x_1, x_2, \dots, x_n)$ and b) $\pi(h|x, x_1, x_2, \dots, x_n)$. To sample from a) we use the fact that this is (almost) the classical kernel so all we need is to generate from an equiprobable mixture of gaussians. To get samples from b) just replace the gaussian kernel into the numerator of equation (14) to obtain, for $\pi(h) \propto h^{-\delta}$,

$$\pi(h|x, x_1, x_2, \dots, x_n) \propto \sum_{\text{all paths}} h^{-(n+\delta+1)} \exp\left\{-\frac{s^2}{2h^2}\right\}. \quad (21)$$

The integral with respect to h of each of the terms being added in (21) is proportional to $s^{-(n+\delta)}$ (see (20)). Thus, by multiplying and dividing by this integral each term, we can write,

$$\pi(h|x, x_1, x_2, \dots, x_n) \propto \sum_{\text{all paths}} \alpha(i) \pi_{s(i)}(h) \quad (22)$$

where $\alpha(i) = (s(i))^{-(n+\delta)}$ as before and,

$$\pi_s(h) \propto h^{-(n+\delta+1)} \exp\left\{-\frac{s^2}{2h^2}\right\} \quad (23)$$

From the change of variables theorem it follows that if y is $\text{Gamma}(\frac{n+\delta}{2}, 1)$ then $h = \sqrt{\frac{s^2}{2y}}$ follows the distribution (23). This shows that the posterior distribution of h is a mixture of transformed gamma distributions. This mixture can be generated by a simple modification to the algorithm used to get the posterior mean of h .

5. Experiments on simulated and real data

I have coded (essentially) the algorithm described in the previous section in MAPLE and tested it dozens of times on simulated data for computing a global value for the smoothness parameter h . All the experiments were carried out with $\delta = 1$

Solid= $0.5N(-1, .5)+0.5N(1, 1)$, $n=50$, 3bumps=CV, 2bumps=MCMC

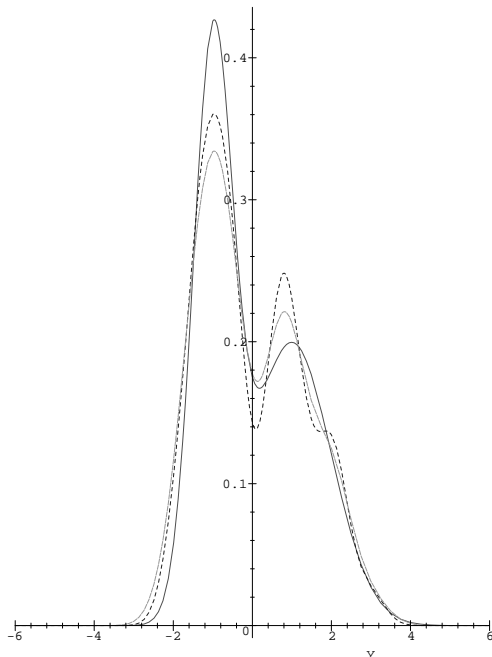


Figure 2. Posterior mean of global h vs plain cross-validation

i.e. with $\pi(h) = h^{-1}$ on mixtures of gaussians. The experiments clearly indicate that the global value of h provided by the MCMC algorithm produce a kernel estimator that is either identical to plain likelihood cross-validation or clearly superior to it depending on the experiment. A typical run is presented in figure [2] where the true density and the two estimators from 50 iid observations are shown. The MAPLE package used in the experiments is available at <http://omega.albany.edu:8008/npde.mpl>.

For comparison with other density estimators in the literature we show in figure [3] the estimate for the complete set of 109 observations of the Old Faithful geyser data. These data are the 107 observations in [2] plus the two outliers 610 and 620. This is a standard gaussian kernel with the global value of $h = 14.217$ chosen by the MCMC algorithm.

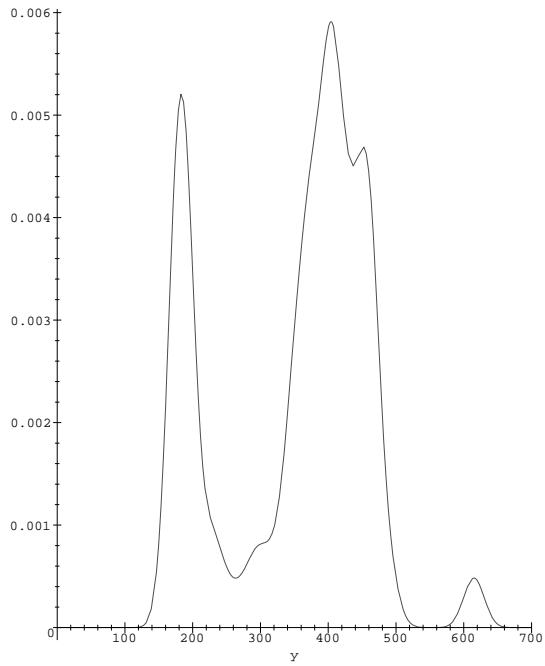


Figure 3. Estimate for the Old Faithful geyser data, $h = 14.217$

6. Conclusions

There is nothing special about dimension one. Only minor cosmetic changes (of the kind: replace h to h^d in some formulas) are needed to include the multivariate case, i.e. the case when the x_j 's are d -dimensional vectors instead of real variables.

Very little is known about these estimators beyond of what it is presented in this paper, In particular nothing is known about rates of convergence. There are many avenues to explore with theory and with simulations but clearly the most interesting and promising open questions are those related to the performance of the QM method above.

References

1. C. C. Rodriguez, "On the estimation of the bandwidth parameter using a non informative prior," in *Proceedings of the 45th session of the ISI*, No. 1, pp. 207-208, International Statistical Institute, August 1985.
2. B. W. Silverman, *Density Estimation: for statistics and data analysis*, Monographs on Statistics and Applied Probability, Chapman and Hall, 1986.