

On the Bayesian Nonparametric Estimation of Discontinuous Densities

Olivier Binette

October 31, 2018

Université du Québec à Montréal

Introduction

The problem:

Estimation of a density f_0 supported on a compact metric space $(M; d)$ using observations X_1, \dots, X_n i.i.d. f_0 .

The problem:

Estimation of a density f_0 supported on a compact metric space $(M; d)$ using observations X_1, \dots, X_n i.i.d. f_0 .

The focus for this talk:

What happens with Bayesian nonparametric procedures when f_0 is discontinuous?

Density f_0 supported on a compact metric space M :

Analysis of non-euclidean data; Statistics on manifolds.

M could be the circle, the sphere, a quotient space, etc.

We want to account for the particular metric structure of M .


Bioinformatics:

Protein structure prediction.



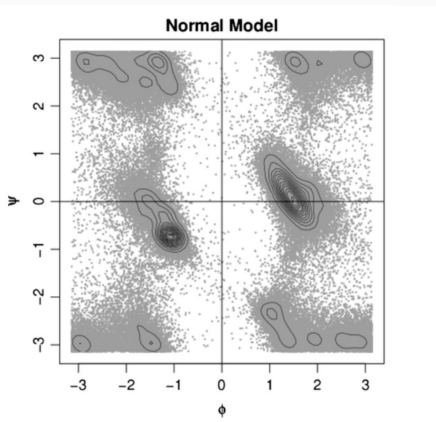
Bioinformatics:

Protein structure prediction.



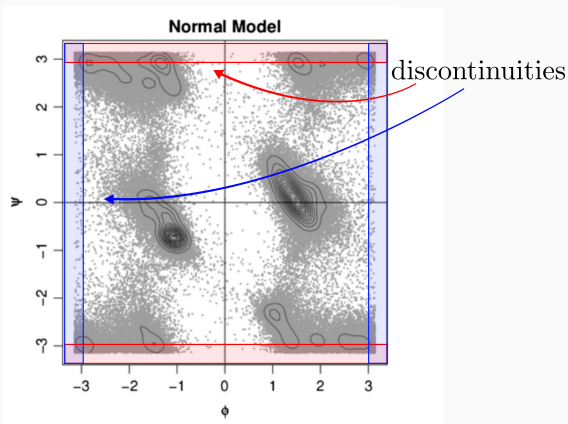
Motivation

Estimation of angle pairs distributions using mixtures of normal distributions (Dhal & al. (2008) *J. Mol. Bio.*).



Motivation

Estimation of angle pairs distributions using mixtures of normal distributions (Dhal & al. (2008) *J. Mol. Bio.*).



Motivation

Failed to take into account the metric structure of the Torus.

Motivation

Failed to take into account the metric structure of the Torus.

The team adjusted their model the next year (Lennox & al. (2009) *JASA*).

Previous work:

Bhattacharya & Dunson (*Strong consistency of nonparametric Bayes density estimation on compact metric spaces with applications to specific manifolds*, 2012) give conditions on DPMs for consistency at positive **continuous** densities.

Previous work:

Bhattacharya & Dunson (*Strong consistency of nonparametric Bayes density estimation on compact metric spaces with applications to specific manifolds*, 2012) give conditions on DPMs for consistency at positive **continuous** densities.

Case $M = [0; 1]$: Petrone and Wasserman (2002) show posterior consistency for random Bernstein polynomial priors at **continuous** densities.

Previous work:

Bhattacharya & Dunson (*Strong consistency of nonparametric Bayes density estimation on compact metric spaces with applications to specific manifolds*, 2012) give conditions on DPMs for consistency at positive **continuous** densities.

Case $M = [0; 1]$: Petrone and Wasserman (2002) show posterior consistency for random Bernstein polynomial priors at **continuous** densities.

Wu and Ghosal (2008) study the Kullback-Leibler property of sieve priors, but also require **continuity**.

Previous work:

Bhattacharya & Dunson (*Strong consistency of nonparametric Bayes density estimation on compact metric spaces with applications to specific manifolds*, 2012) give conditions on DPMs for consistency at positive **continuous** densities.

Case $M = [0; 1]$: Petrone and Wasserman (2002) show posterior consistency for random Bernstein polynomial priors at **continuous** densities.

Wu and Ghosal (2008) study the Kullback-Leibler property of sieve priors, but also require **continuity**.

Contrasts with KDE which are known to be **L^1 consistent at bounded discontinuous densities** since Devroye & Wagner (1979).

What happens when we drop continuity assumptions on f_0 ?

Our contribution

Show that posterior consistency also holds at discontinuous (bounded) densities for a class of sieve priors.

Our contribution

Show that posterior consistency also holds at discontinuous (bounded) densities for a class of sieve priors.

In particular, we show that continuity assumptions are not necessary for the random (Bernstein) polynomial priors of Petrone and Wasserman (2002).

Background

\mathcal{F} as set of densities; a prior on \mathcal{F} .

Kullback-Leibler divergence between f_0 and f :

$$K(f_0; f) = \int_{f_0 > 0} f_0 \log \frac{f_0}{f}:$$

\mathcal{F} as set of densities; a prior π on \mathcal{F} .

Kullback-Leibler divergence between f_0 and f :

$$K(f_0; f) = \int_{f_0 > 0} f_0 \log \frac{f_0}{f}$$

Why is it important? If X_i is i.i.d. f_0 , then

$$K(f_0; f) = L \left(\prod_{i=1}^k \frac{f_0(X_i)}{f(X_i)} = e^{Lk + o(k)} \right) \text{ a.s. as } k \rightarrow \infty$$

The **Kullback-Leibler support** of f_0 is the set of all densities $f \in \mathcal{F}$ such that

$$\{ f \in \mathcal{F} : K(f_0; f) < \epsilon \} > 0 \quad \forall \epsilon > 0:$$

The **Kullback-Leibler support** of \mathcal{F} is the set of all densities $f \in \mathcal{F}$ such that

$$\left(f \in \mathcal{F} : K(f_0; f) < \epsilon \right) > 0 \quad \forall \epsilon > 0:$$

The posterior distribution of θ is strongly consistent at every density of its Kullback-Leibler support provided a complexity constraint on the prior (fundamental work by Andrew Barron and many others; see also Xing & Ranneby (2009)).

Strong posterior consistency means that if X_i is i.i.d. f_0 , then

$$\int |f - f_0| dP_{X_i}^k \rightarrow 0 \quad \text{a.s. as } k \rightarrow \infty$$

a.s. as $k \rightarrow \infty$ and for every $\epsilon > 0$.

Implies almost sure convergence of the Bayes estimator.

The Kullback-Leibler support of sieve priors

$(M; d)$ a compact metric space together with a finite measure defined on its Borel σ -algebra.

General framework

$(M; d)$ a compact metric space together with a finite measure defined on its Borel σ -algebra.

F a set of **bounded** densities.

General framework

$(M; d)$ a compact metric space together with a finite measure defined on its Borel σ -algebra.

F a set of **bounded** densities.

Linear operators $T_n : L^1(M) \rightarrow L^1(M)$, $n \in \mathbb{N}$ which **maps densities to densities** and $T_n(F) \subseteq F$.

General framework

$(M; d)$ a compact metric space together with a finite measure defined on its Borel σ -algebra.

F a set of **bounded** densities.

Linear operators $T_n : L^1(M) \rightarrow L^1(M)$, $n \in \mathbb{N}$ which **maps densities to densities** and $T_n(F) \subseteq F$.

We thus obtain a model $C = \sum_{n \in \mathbb{N}} C_n$ with $C_n := T_n(F)$.

General framework

$(M; d)$ a compact metric space together with a finite measure defined on its Borel σ -algebra.

F a set of **bounded** densities.

Linear operators $T_n : L^1(M) \rightarrow L^1(M)$, $n \in \mathbb{N}$ which **maps densities to densities** and $T_n(F) \subseteq F$.

We thus obtain a model $C = \sum_{n \in \mathbb{N}} C_n$ with $C_n := T_n(F)$. Priors π_n on C_n and distribution π on \mathbb{N} yields

$$= \sum_n \pi(n) \pi_n:$$

General framework

T_n is linear, maps densities to densities and $T_n(F) \in \mathcal{P}(F)$.
 $\int T_n f = \int f$.

Theorem (Binette & Guilote, 2018)

Let F , $\mathcal{P}(F)$, and T_n be as before. Suppose also that

$T_n(F)$ is of finite dimension, and

$\|T_n f\|_1 \leq k_1 \|f\|_1$ for every continuous f .

If $k_1 > 0$ and $\mathcal{P}(F)$ has L^1 support $T_n(F)$, then any (bounded) $f_0 \in \mathcal{P}(F)$ is in the Kullback-Leibler support of T_n .

Corollary

Under the hypotheses of the Theorem, suppose further that $\epsilon_n < C e^{-c d_n}$ for some increasing sequence $d_n > \dim T_n(F)$ and some $c, C > 0$. Then the posterior distribution of μ is strongly consistent at any (bounded) $f_0 \in F$.

Corollary

Under the hypotheses of the Theorem, suppose further that $\epsilon_n < C e^{-c d_n}$ for some increasing sequence $d_n > \dim T_n(F)$ and some $c, C > 0$. Then the posterior distribution of μ is strongly consistent at any (bounded) $f_0 \in F$.

Idea: If $\epsilon_n < C e^{-c d_n}$, then the prior satisfies a finite Hausdorff ϵ_n -entropy conditions which entails posterior consistency on the Kullback-Leibler support (Xing & Ranneby (2009); same spirit as Walker (2004) and Barron (1986)).

What can this be applied to?

Sieve priors.

Random series priors / finite mixtures with a random number of components.

Dirichlet Process location mixtures with a discrete scale parameter;

$$T_n f = f_{1=n}.$$

Random polynomial priors

The Bernstein-Kantorovich operator is defined as

$$T_n f(x) = (n+1) \int_0^1 f(t) dt p_{i;n}(x); \quad x \in [0; 1];$$

where $p_{i;n}(u) = \binom{n}{i} u^i (1-u)^{n-i}$ is the i th Bernstein polynomial of degree n .

The Bernstein-Kantorovich operator is defined as

$$T_n f(x) = (n+1) \int_0^1 f(t) dt p_{i;n}(x); \quad x \in [0; 1];$$

where $p_{i;n}(u) = \binom{n}{i} u^i (1-u)^{n-i}$ is the i th Bernstein polynomial of degree n . It may be extended to act on a probability measure P by letting

$$T_n P = (n+1) \int_0^1 P \left[\frac{i}{n+1}; \frac{i+1}{n+1} \right] p_{i;n}(x):$$

If P has a density f , then $T_n P = T_n f$.

Bernstein-Dirichlet prior of Petrone and Wasserman (2002): DPM induced by the random density

$T_N D$; where N ; D DP:

Bernstein-Dirichlet prior of Petrone and Wasserman (2002): DPM induced by the random density

$$T_N D; \text{ where } N \text{ ; } D \text{ DP:}$$

T_n is a linear operator mapping densities to densities; its image is of dimension n and it is well-known that $\|T_n f - f\|_1 \leq 0$ if $f \in C([0; 1])$.

Bernstein-Dirichlet prior of Petrone and Wasserman (2002): DPM induced by the random density

$$T_N D; \text{ where } N \rightarrow \infty; D \sim \text{DP}:$$

T_n is a linear operator mapping densities to densities; its image is of dimension n and it is well-known that $\|T_n f - f\|_1 \leq C \|f\|_2$ for $f \in C([0; 1])$.

Provided $T_n(D)$ has full support and $\|T_n f - f\|_1 \leq C \|f\|_2$, then posterior consistency holds at any bounded

Random polynomial priors

Bernstein-Dirichlet prior of Petrone and Wasserman (2002): DPM induced by the random density

$$T_N D; \text{ where } N \rightarrow \infty; D \text{ DP:}$$

T_n is a linear operator mapping densities to densities; its image is of dimension n and it is well-known that $\|T_n f - f\|_1 \leq C \|f - T_n f\|_1$ for $f \in \mathcal{C}([0;1])$.

Provided $T_n(D)$ has full support and $0 < \inf_n \int T_n(D) < C e^{-cn}$, then posterior consistency holds at any bounded f_0 .

This shows the continuity assumptions of Petrone and Wasserman (2002) are not necessary; **boundedness is sufficient**.

Conclusion

Conclusion

For a class of sieve priors, L^1 posterior consistency also holds at bounded discontinuous densities.

- Circular/spherical density estimation
- Copula density estimation and other shape constrained problems.

Notes and slides can be found on my personal page
olivierbinette.ca.

Full story is in *Bayesian Nonparametrics for Directional Statistics* (Binette & Guillotte, 2018; arXiv:1807.00305).

Thank you!