Published in final edited form as:

J Comput Graph Stat. 2016 ; 25(1): 301–320. doi:10.1080/10618600.2014.1002927.

Computational Aspects of Optional Pólya Tree

Hui Jiang,

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109

John Chong Mu,

Department of Electrical Engineering, Stanford University, Stanford, CA 94305

Kun Yang,

Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305

Chao Du,

Department of Statistics, University of Virginia, Charlottesville, VA 22904

Luo Lu, and

Twitter Inc., San Francisco, CA 94103

Wing Hung Wong

Department of Statistics and Health Research and Policy, Stanford University, Stanford, CA 94305

John Chong Mu: johnmu@umhoj.com; Kun Yang: kunyang@stanford.edu; Chao Du: cd2wb@virginia.edu; Luo Lu: llu@twitter.com

Abstract

Optional Pólya tree (OPT) is a flexible nonparametric Bayesian prior for density estimation. Despite its merits, the computation for OPT inference is challenging. In this paper we present time complexity analysis for OPT inference and propose two algorithmic improvements. The first improvement, named limited-lookahead optional Pólya tree (LL-OPT), aims at accelerating the computation for OPT inference. The second improvement modifies the output of OPT or LL-OPT and produces a continuous piecewise linear density estimate. We demonstrate the performance of these two improvements using simulated and real data examples.

Keywords

density estimation; Bayesian nonparametrics; recursive partition; smoothing; time complexity

1 Introduction

Given independent random samples, characterizing their distribution is one of the most fundamental problems in statistics. Density estimation is a key approach to solving such

problems and is closely related to other problems such as clustering, classification and regression [Silverman, 1986]. Kernel density estimation [Rosenblatt, 1956, Parzen, 1962] and data-driven histograms [Scott, 1979, Barron et al., 1992] are popular methods for density estimation. Among many approaches developed for density estimation, nonparametric Bayesian models have the merits of being highly flexible and computationally feasible. In a seminal paper, Ferguson [1973] proposed a class of prior distributions named Dirichlet process priors which have flexible structures and analytically tractable posterior distributions which make them very popular for nonparametric Bayesian modeling. However, a random probability measure sampled from a Dirichlet process prior is discrete with probability one, which makes it unsuitable for density estimation. Using the Dirichlet process to model an infinite mixture of distributions gives rise to absolutely continuous distributions [Lo, 1984]. In practice, Dirichlet process Gaussian mixture models (DP-GMM) is commonly used for density estimation [Escobar and West, 1995]. Detailed discussions of Dirichlet process and its extensions as well as other Bayesian nonparametric models can be found in Ghosh and Ramamoorthi [2003].

Proposed by Ferguson [1974] and studied in details by Lavine [1992], Pólya tree priors are extensions of Dirichlet processes [Ferguson, 1974] and special cases of tailfree processes [Freedman, 1963]. A Pólya tree defines a prior distribution which gives rise to random probability measures on a separable measurable space $(\Omega, \mathcal{B}(\Omega))$, with respect to a base measure μ . In this paper we assume that Ω is a bounded hyperrectangle in $\mathbb{R}^p$, $\mathcal{B}(\Omega)$ is the Borel σ -algebra on Ω and μ is the Lebesgue measure. The theory applies to more general cases, e.g., Ω is a finite set and μ is the counting measure.

The Pólya tree prior [Lavine, 1992] works as follows: starting from the root region $A = \Omega$, it defines a partition of A as $A = A_0 \cup A_1$, where A_0 and A_1 are disjoint. It then defines partitions for A_0 and A_1 as $A_0 = A_{00} \cup A_{01}$ and $A_1 = A_{10} \cup A_{11}$, respectively. This partition process recursively goes on forever: for any binary sequence $s = s_1 \dots s_k$ where $s_i \in \{0, 1\}$, $i = 1, \dots, k$, a partition of the subregion A_s is defined as $A_s = A_{s0} \cup A_{s1}$ with A_{s0} and A_{s1} disjoint. A random probability measure Q can be generated from the Pólya tree prior as follows: for each region A_s , where $s = s_1 \dots s_k$ is a binary sequence and let $A_s = A$ when $s = \varepsilon$ is the empty sequence (i.e., $k = 0$), a random variable Y_s is sampled from a Beta distribution with parameters α_{s0} and α_{s1} , where α_{s0} and α_{s1} are parameters of the Pólya tree. Finally, let

$$Q(A_s) = \left(\prod_{j=1; s_j=0}^k Y_{s_1 \dots s_{j-1}} \right) \left(\prod_{j=1; s_j=1}^k (1 - Y_{s_1 \dots s_{j-1}}) \right)$$

for any subregion A_s , and let $Q(A) = 1$.

In order for the Pólya tree prior to produce absolutely continuous measures, the Beta parameters α_s have to increase rapidly as the depth of partitioning increases [Kraft, 1964, Mauldin et al., 1992]. Furthermore, because the partitioning process goes on forever, with probability one the density sampled from a Pólya tree prior will have discontinuity almost

everywhere. To solve these problems, Wong and Ma [2010] proposed optional Pólya tree (OPT), which is an extension to the Pólya tree prior that allows for optional stopping and random selection of partitioning variables. It has three attractive properties: 1) It fits the data adaptively; 2) It produces absolutely continuous distributions; and 3) The computation for exact OPT inference can be done in finite steps, providing there are no replicates in the data. OPT has been successfully applied in both one-sample [Wong and Ma, 2010] and two-sample [Ma and Wong, 2011] settings.

Albeit all these merits, there are existing issues which may hinder the application of OPT. Two major issues are: 1) The running time for exact OPT inference increases rapidly as sample size and the dimension of the data increase, which becomes prohibitive when the dimension is high; 2) The posterior mean of OPT is usually computationally prohibitive to estimate [Wong and Ma, 2010], and although the hierarchical maximum a posteriori (hMAP) estimate of OPT proposed by Wong and Ma [2010] is computationally more friendly, it gives rise to piecewise constant density estimates where discontinuities occur at boundaries between the pieces which could become undesirable for applications where continuous densities are sought.

In this paper we attempt to tackle the above two problems. In particular, we demonstrate that the computation for OPT inference can be greatly accelerated by performing approximate inference. We also present an approach which modifies the output of OPT and produces a continuous piecewise linear density estimate. This paper is organized as follows: In Section 2 we review the construction of OPT and the algorithm for its exact inference. In Section 3 we present an approach for reducing the computation named limited-lookahead optional Pólya tree (LL-OPT). Section 4 gives the time complexity analysis for OPT and LL-OPT. In Section 5 we present an approach for the estimation of smooth densities. Simulated and real data examples are given in Section 6. Further discussions and software availability are given in Section 7.

A personal computer with Intel Xeon X7560 2.27Ghz on a single core was used for all computations reported in this paper.

2 Optional Pólya Tree and its inference

Similar to Pólya trees, OPT operates simultaneously as two processes: the first process recursively and randomly partitions the space Ω and the second process recursively assigns a random probability measure into each region in the partition generated by the first process.

The recursive partitioning process works as follows: starting from the root region $A = \Omega$, a random binary sample S is drawn according to $\text{Bernoulli}(\rho(A))$ where $\rho(A)$ is a parameter of the OPT. If $S = 1$, we stop further partitioning A and mark it as a terminal region. If $S = 0$, a random integer $J \in \{1, \dots, J_A\}$ is drawn according to a probability vector $\lambda(A) = (\lambda_1(A), \dots, \lambda_{J_A}(A))$, where J_A is the number of different ways of partitioning A into disjoint subregions

and $(\lambda_1(A), \dots, \lambda_{J_A}(A))$ is a set of parameters of the OPT such that $\sum_{j=1}^{J_A} \lambda_j(A) = 1$. If $J = j$,

A is partitioned in the j -th way into I_A^j disjoint subregions $\{A_{I_A^j}^1, \dots, A_{I_A^j}^{I_A^j}\}$ such that

$$A = \bigcup_{i=1}^{I_A^j} A_i^j, \\ \forall i, k \in \{1, \dots, I_A^j\}, \text{ if } i \neq k, \text{ then } A_i^j \cap A_k^j = \emptyset.$$

The partitioning process is then recursively applied to each subregion A_i^j of A , and consequently, a partition tree is built by the process recursively. The partitioning process terminates when all the regions in the partition are marked as terminal regions.

The process for random assignment of probability measures works along with the partitioning process. It begins with a probability measure $Q^{(0)}$ which is uniform within the root region $A = \Omega$. Whenever a region A is partitioned into subregions $\{A_i^j\}_{i=1}^{I_A^j}$, a random vector $\theta_A^j = (\theta_{A,1}^j, \dots, \theta_{A,I_A^j}^j)$ is drawn from a Dirichlet distribution with parameter $\alpha^j(A) = (\alpha_1^j(A), \dots, \alpha_{I_A^j}^j(A))$ and then a new probability measure $Q^{(k+1)}$ is built based on $Q^{(k)}$ by reassigning the probability mass within A according to $Q^{(k+1)}(A_i^j) = \theta_{A,i}^j Q^{(k)}(A)$, where $k = k(A)$ is the level of region A in the partition tree with $k(\Omega) = 0$. Within each subregion A_i^j , the probability mass is again uniformly distributed.

For the OPT procedure described above, there exists a limiting probability measure for $Q^{(k)}$ [Wong and Ma, 2010, Theorem 1]: If $\rho(A)$ is uniformly bounded away from 0 and 1 for all A , then $Q^{(k)}$ converges almost surely in variational distance to a probability measure Q which is absolutely continuous with respect to μ . The density of Q with respect to μ on Ω is piecewise constant with countably many pieces.

As a prior distribution on the space of probability measures, one merit of OPT is that the computation for exact OPT inference is analytically manageable. Given independent random samples $\mathbf{x} = (x_1, \dots, x_n)$ from Q when Q has a OPT prior π with independent stopping probabilities $\rho(A)$ and all the parameters of π are uniformly bounded away from 0 and 1, the posterior distribution of Q is also an OPT with the following parameters [Wong and Ma, 2010, Theorem 3]:

1. Stopping probabilities: $\rho(A|\mathbf{x}) = \rho(A)\Phi_0(A)/\Phi(A)$,
2. Selection probabilities:

$$\lambda_j(A|\mathbf{x}) \propto \lambda_j(A) \frac{D(\mathbf{n}_A^j + \alpha^j(A))}{D(\alpha^j(A))} \prod_{i=1}^{I_A^j} \Phi(A_i^j), j=1, \dots, J_A,$$

3. Probability mass allocation:

$$\alpha^j(A|\mathbf{x}) = \alpha^j(A) + \mathbf{n}_A^j, j=1, \dots, J_A,$$

where $\mathbf{n}_A^j = (n_{A,1}^j, \dots, n_{A,I_A^j}^j)$ are the numbers of samples in $\mathbf{x}$ falling into each regions of the partition, $D(\mathbf{t}) = \Gamma(t_1) \cdots \Gamma(t_k) / \Gamma(t_1 + \cdots + t_k)$, $\Phi_0(A) = \mu(A)^{-n(A)}$ where $n(A)$ is the number of samples in region A , and $\Phi(A)$ can be calculated using the following recursive equation until a terminal region with either 0 or 1 data point is reached:

$$\Phi(A) = \rho(A)\Phi_0(A) + (1 - \rho(A)) \sum_{j=1}^{J_A} \lambda_j(A) \frac{D(\mathbf{n}_A^j + \boldsymbol{\alpha}^j(A))}{D(\boldsymbol{\alpha}^j(A))} \prod_{i=1}^{I_A^j} \Phi(A_i^j). \quad (2.1)$$

There are many different schemes for partitioning the regions. In this paper we consider a simple binary partitioning scheme proposed in Wong and Ma [2010]: For $A = \{(t_1, \dots, t_p) : t_j \in [l_j, u_j]\}$, i.e., a hyperrectangle in $\mathbb{R}^p$, there are $J_A = p$ ways to partition A . For the j -th way, $j = 1, \dots, p$, we partition A into two subregions at the midpoint of the range of t_j such that $A_1^j = \{t \in A : t_j < (l_j + u_j)/2\}$ and $A_2^j = A \setminus A_1^j$.

3 Acceleration of computation

According to (2.1), exact OPT inference requires computing $\Phi(A)$ recursively, which is computationally challenging since the running time increases rapidly as the dimension of the data increases. This is evident based on the fact that in (2.1) the value of $\Phi(A)$ has to be computed for every region A with at least one data point and the number of such regions is enormous when the dimension is high. Figure 1 shows the running time for OPT inference for samples of various sizes and dimensions, using a fast algorithm for exact OPT inference which will be introduced Section 4. We can see that the running time increases nearly exponentially with respect to the dimension of the data and it becomes prohibitive when the dimension is 8 or higher. The analysis of the time complexity for OPT inference is given in Section 4.

The computation for OPT inference can be greatly accelerated by performing the inference approximately. In a straightforward approach, named naive inexact optional Pólya tree (NI-OPT), we stop the recursion in (2.1) when the number of samples in a region A is small or when the volume of A is small, and compute $\Phi(A)$ by assuming that all the samples in A are uniformly distributed and therefore no further partitioning of A is necessary. The rationale for employing these heuristics is that when the volume of A is small, error in $\Phi(A)$ is only likely to affect the estimated density locally, and when the number of samples in A is small, $\Phi(A)$ would not be too different from $\Phi'(A)$ which is computed by redistributing the samples in A uniformly.

In practice, we found that the running time for the NI-OPT approach, although significantly reduced compared with the exact OPT approach, is still prohibitive when the sample size is large and when the dimension is high. In this paper we propose a more efficient algorithm for approximate OPT inference, named limited-lookahead optional Pólya tree (LL-OPT), which works as follows: starting from the root region $A = \Omega$, we perform the recursive computation in (2.1) for up to h levels in the partition tree, where h is a tuning parameter which provides a trade-off between running time and estimation accuracy. For a region B at

level $k(B) = k(A) + h$, we stop the recursion and approximate $\Phi(B)$ as $\Phi'(B)$ using the uniform approximation as we did in NI-OPT. Based on the computed $\Phi(A)$ and $\Phi(A_s)$ where $A_s (s \in S)$ are all the subregions of A , we compute the hierarchical maximum a posteriori (hMAP) tree using the approach described in Wong and Ma [2010] for up to q levels, where $q - h$ is another tuning parameter. In our implementation we fix $q = 1$, as this is the least aggressive choice and we found it produces sufficiently efficient computation. For each leaf node in the hMAP tree, we recursively apply the LL-OPT algorithm and build a sub-tree by treating that leaf node as the root node. Furthermore, like in NI-OPT, we can also stop the procedure when the number of samples in the region is small enough, or when the volume of the region is small enough.

There are some connections to existing methods for Pólya tree inference. Lavine [1992] studied scenarios where exact inference for Pólya tree can be achieved, and Mauldin et al. [1992], Lavine [1994] studied scenarios where the inference for Pólya tree is only performed upto a given level under which errors of approximation can be estimated or bounded. Mixture of Pólya tree priors (MPT) is another approach for estimating smooth densities [Lavine, 1992]. The idea of early stopping was first proposed by Hutter [2009]. Bayesian recursive partitioning has also been proposed previously as a classification problem [Denison et al., 1998].

4 Time complexity analysis for OPT and LL-OPT

In this section we study the time complexity for the exact OPT and the LL-OPT algorithms. For simplicity we consider only the simple binary partitioning scheme described in Section 2. However, the analysis technique may be adapted to more complicated partitioning schemes. For each region A , the exact OPT algorithm according to (2.1) has to undertake the following three computational tasks:

1. For each subregion A_i^j of A , $j = 1, \dots, p$, $i = 1, 2$, determine $n_{A,i}^j$, the counts of samples within A_i^j .
2. For each subregion A_i^j of A , compute $\Phi(A_i^j)$ recursively using (2.1).
3. Calculate $\Phi(A)$ according to (2.1) based on the counts found in task 1 and $\Phi(A_i^j)$ computed for each subregion A_i^j in task 2.

For task 1, as $n \rightarrow \infty$, the straightforward way of counting the number of samples in each of the $2p$ subregions of A will take $O(2pN)$ operations, where N is the total number of samples in Ω . However, this cost can be reduced by trading space for time. If we can store the set of samples in A , counting the number of samples in each of the $2p$ subregions of A will only take $O(2pn)$ operations, where n is the number of samples in A . Let $f(n)$ denote the time for computing $\Phi(A)$ for a region A with n samples. The time complexity for task 2 is therefore $\sum_{j=1}^p \{f(n_{A,1}^j) + f(n_{A,2}^j)\}$. For task 3, the time complexity is $O(2p)$ if we assume that function $D(t)$ can be computed in $O(1)$ time. The time complexity for computing $\Phi(A)$ is therefore

$$\begin{aligned}
 f(n) &= O(2pn) + \sum_{j=1}^p \{f(n_{A,1}^j) + f(n_{A,2}^j)\} + O(2p) \\
 &= O(pn) + \sum_{j=1}^p \{f(n_{A,1}^j) + f(n_{A,2}^j)\}
 \end{aligned} \quad (4.1)$$

It is shown in (4.1) that the value of $f(n)$ depends on the values of $n_{A,1}^j$ and $n_{A,2}^j$, which are determined by the distribution of samples within A . The more uniformly the samples are distributed in A , the less the computation time is required. Here we estimate $f(n)$ for a range of scenarios.

Theorem 4.1. For the following three cases, the time complexities for the exact OPT algorithm are $f(n) = O(pn^{\log 1/2^p})$ (case 1), $f(n) = O(n^{2^p-1})$ (case 2) and $f(n) = O(p^{n/c})$ (case 3), respectively, where n is the sample size and p is the dimension.

1. $n_{A,1}^j = tn$ and $n_{A,2}^j = (1-t)n$, where $0.5 < t < 1$ is a constant.
2. $n_{A,1}^j = Tn$ and $n_{A,2}^j = (1-T)n$, where T is uniformly distributed within $(0, 1)$.
3. $n_{A,1}^j = c$ and $n_{A,2}^j = n - c$, where $c > 0$ is a constant.

We see that case 2 represents a more realistic picture between the two extremes in case 1, namely $t = 0.5$ and t is close to 1, and case 3 is the limit of case 1 as $1-t$ goes to c/n . In effect, case 1 indicates that the number of levels to be explored is $\log_{1/t} n$, and case 3 indicates that the number of level is n/c . Although only 3 specific cases are discussed here, the results actually provide full coverage over general scenarios.

The simplest implementation of the OPT algorithm is to use depth-first search, which we call DF-OPT. We find DF-OPT to be very inefficient in practice, partly because regions that can be reached through many different paths of binary partitions will have their $\Phi(\cdot)$ values computed repeatedly, which is unnecessary. A more efficient implementation is to cache all the computed $\Phi(\cdot)$ values so that for each region A , $\Phi(A)$ will need to be computed only once. In practice we find that this cached implementation can accelerate computation substantially. Therefore, we use the cached implementation as the default implementation for all the simulations (denoted as OPT in experiments). Comparisons between DF-OPT and cached OPT are given in Section 6. It is worthwhile pointing out here that although cached OPT is much faster than DF-OPT, it only works with limited sample sizes and dimensions because it uses much larger amount of memory. In contrast, DF-OPT has the advantage that its memory usage is minimal. Furthermore, if a partitioning scheme other than the binary partitioning is used, caching might become less effective or even non-effective.

We found that in practice, the cached implementation of OPT is more likely to be limited by memory size rather than by running time. Therefore, we also study the space complexity of the cached OPT here.

Proposition 4.2. Under the binary partitioning scheme, the total number of regions at level k is $\binom{k+p-1}{k} 2^k$, where p is the dimension of the sample space. The total number of regions up to level k is therefore $\sum_{i=0}^k \binom{i+p-1}{i} 2^i$.

If we assume that the depth of the partition tree is on the order of $\log_2 n$ where n is the sample size, which is the case when the samples in Ω are distributed roughly uniformly, the

space required by the cached OPT is $O\{\sum_{i=0}^{\log_2 n} (i+p-1)2^i\} = O\{n((\log_2 n)+p-1)\}$.

The LL-OPT algorithm is more friendly in CPU and memory requirements because for each cut it only looks ahead for h steps, and as a result, the time and space complexity for a single cut is bounded.

Theorem 4.3. For the LL-OPT algorithm with parameter h , the time complexity is $f(n) = O\{n(n + 2^h p)p^h\}$, where n is the sample size and p is the dimension.

Theorems 4.1 and 4.3 show that the reduction on running time using the LL-OPT algorithm can be significant, especially when n and p are large. The amount of reduction is controlled by the tuning parameter h . The smaller the h , the faster the LL-OPT algorithm. However, a smaller h also leads to less accurate estimation. The following adaptive approach can be used for selecting h : run the LL-OPT algorithm many times starting with $h = 1$ and increase h by 1 each time until a stopping decision is made (either based on a cost constraint or a criterion evaluating the gain between two successive estimates). Theorem 4.3 shows that the total time for running the LL-OPT algorithm for $h = 1, \dots, k-1$ is less than the run time for a single $h = k$.

5 Estimation of smooth density

The densities estimated by OPT are piecewise constant and have discontinuities at boundaries between the pieces, which could become undesirable for applications where continuous densities are sought. In this section we introduce a procedure which constructs a continuous piecewise linear approximation to the piecewise constant density estimated by OPT via quadratic programming. We call the continuous piecewise linear density estimator constructed using this approach the finite element estimator (FEE), because it was inspired by the idea of domain partitioning in finite element method [Allaire, 2007].

We define a proper space for continuous piecewise linear densities as follows. We start with the definition of a triangulation of the domain Ω (without loss of generality, we shall assume Ω is $[0, 1]^p$) by p -simplices. A p -simplex in $\mathbb{R}^p$ is the convex envelop of its $p+1$ vertices $\{a_j\}_{j=1}^{p+1}$, which is nondegenerate if its vertices do not fall on a hyperplane in $\mathbb{R}^p$ and consequently the volume of Δ , denoted as $\mu(\Delta)$, is not zero, which we shall always assume in what follows.

Definition 5.1. Let $\Omega = [0, 1]^p$, a triangulation of Ω is a set Γ of p -simplices $\{\Delta_i\}_{i=1}^m$ satisfying

- a. $\Delta_i \subset \Omega$ ($i = 1 \dots m$) and $\Omega = \bigcup_{i=1}^m \Delta_i$
- b. For any two distinct simplices Δ_i and Δ_j if they share k ($1 \leq k \leq p$) vertices, $\Delta_i \cap \Delta_j$ is a $(k-1)$ -simplex formed by these k vertices, or an empty set otherwise. This condition guarantees that no vertex of Δ_i lies on the faces of Δ_j and vice versa. This can help exclude the undesired situations as shown in Figure 2.

Definition 5.2. For an OPT partition $\Omega = \bigcup_{i=1}^l A_i$, a triangulation of the OPT partition is a triangulation Γ of Ω where each region A_i is also triangulated by a subset of Γ . See Figure 3 for an example.

Definition 5.3. Given a triangulation Γ of an OPT partition of Ω , a function $f: \Omega \rightarrow \mathbb{R}$ is called continuous piecewise linear over Γ if f is linear within each simplex in Γ . The set of continuous piecewise linear functions over Γ is denoted as $\mathcal{P}_1(\Gamma)$.

Let the set $\{a_1 \dots a_k\}$ be all the vertices in Γ , we define k continuous piecewise linear functions as $\varphi_i(a_j) = \delta_{ij}$ for all $1 \leq i, j \leq k$. In another word, φ_i is 1 at a_i and 0 at all other vertices. The set of functions $\varphi_i (i = 1 \dots k)$ forms a basis for all continuous piecewise linear functions over Γ [Allaire, 2007]. One such basis function is shown in Figure 4.

It has been shown that the piecewise constant density estimated by OPT is weakly consistent [Wong and Ma, 2010, Theorem 4]. As the sample size grows, all the probability will be asymptotically concentrated in any weak neighborhood of the true distribution with bounded density. Here, we aim at achieving the following two properties with the continuous piecewise linear FEE.

1. (fidelity property) The difference between the piecewise constant density and the continuous piecewise linear density must be small.
2. (smoothness property) The variation of the continuous piecewise linear density within each region of the OPT partition must be small.

Unfortunately, in practice, these two properties often contradict each other. Therefore, we use penalized optimization to achieve a trade-off. Let $f \in \mathcal{P}_1(\Gamma)$ be a density function which is continuous piecewise linear over Γ , the fidelity property can be enforced by the following penalty $p(\cdot)$

$$p(f) = \sum_{j=1}^m w_j \left[\int_{\Delta_j} f(x) dx - Q(\Delta_j) \right]^2,$$

where Δ_j is the j -th simplex in Γ , $Q(\Delta_j)$ is the probability mass assigned in Δ_j by the piecewise constant OPT density estimate and w_j is a pre-specified weight depending on the importance of simplex j . In our simulations, we find that choosing w_j to be $\mu(\Delta_j)^{-2}$ provides good results by keeping a balance between large and small simplices.

The smoothness property implies that the density function in any given simplex should be as flat as possible. Therefore, in a density plot (see Figure 5 for an example), the volume of the simplex on the top, denoted as $\mu^*(\Delta_j)$, should be approximately equal to the volume of the simplex on the bottom, which is $\mu(\Delta_j)$. Therefore, the smoothness property can be enforced by the following penalty $q(\cdot)$

$$q(f) = \sum_{i=1}^m \left[\frac{\mu * (\Delta_j)}{\mu(\Delta_j)} \right]^2.$$

To enforce both properties simultaneously, we minimize $h(f) = p(f) + \lambda q(f)$, where λ is a tuning parameter. Therefore, the continuous piecewise linear density estimate f can be obtained by solving the following optimization problem.

$$\begin{aligned} f = \arg \min_{f \in \mathcal{P}_1(\Gamma)} & \sum_{j=1}^m w_j \left[\int_{\Delta_j} f(x) dx - Q(\Delta_j) \right]^2 \\ & + \lambda \sum_{i=1}^m \left[\frac{\mu * (\Delta_j)}{\mu(\Delta_j)} \right]^2, \text{ subject to } f(x) \geq 0 \text{ for all } x \in \Omega \text{ and } \sum_{i=1}^m \int_{\Delta_i} f(x) dx \\ & = \int_{\Omega} f(x) dx = 1. \end{aligned} \quad (5.1)$$

Within a given simplex Δ_j with vertices $\{a_1, \dots, a_{p+1}\}$, and with $\{c_1, \dots, c_{p+1}\} = \{f(a_1), \dots, f(a_{p+1})\}$ as the corresponding densities on the vertices, $f(x)$ can be written as a linear combination of the basis functions

$$f(x) = \sum_{i=1}^{p+1} c_i \phi_i(x),$$

where $c_i \geq 0$. We have

$$\mu(\Delta_j) = \frac{1}{p!} \left| \det \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,p} & 1 \\ a_{2,1} & a_{2,2} & \cdots & a_{2,p} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{p+1,1} & a_{p+1,2} & \cdots & a_{p+1,p} & 1 \end{pmatrix} \right|, \int_{\Delta_j} f(x) dx = \frac{\mu(\Delta_j)}{p+1} (c_1 + \cdots + c_{p+1}).$$

and

$$\mu * (\Delta_j) = \frac{1}{p!} \left\| \det \begin{pmatrix} \mathbf{i}_1 & \cdots & \mathbf{i}_p & \mathbf{i}_{p+1} \\ a_{2,1} - a_{1,1} & \cdots & a_{2,p} - a_{1,p} & c_2 - c_1 \\ \vdots & \vdots & \vdots & \vdots \\ a_{p+1,1} - a_{1,1} & \cdots & a_{p+1,p} - a_{1,p} & c_{p+1} - c_1 \end{pmatrix} \right\|_2$$

where $\|\cdot\|_2$ is the Euclidean norm and $\mathbf{i}_j (j = 1, \dots, p+1)$ are the natural basis of $\mathbb{R}^{p+1}$. Since $\int_{\Delta_j} f(x) dx$ is a linear function of $c_1, \dots, c_n$ and $\mu * (\Delta_j)^2$ is a quadratic function of $c_1, \dots, c_n$, (5.1) can be solved by quadratic programming.

6 Examples

6.1 LL-OPT

We compare the running times and the estimation accuracies of the exact OPT and the LL-OPT algorithms. In our simulations, the root region Ω is the unit cube $[0, 1]^p$ in $\mathbb{R}^p$. The exact OPT algorithm and the LL-OPT algorithm are used to estimate the density based on the simulated samples with different sizes. The Hellinger distance

$H(f, g) = \left\{ \frac{1}{2} \int (f(x)^{1/2} - g(x)^{1/2})^2 dx \right\}^{1/2}$ was used as the metric to evaluate the accuracy of density estimates.

Example 6.1. We consider a mixture distribution of two independent components over the unit square $[0, 1] \times [0, 1]$ [Wong and Ma, 2010, Example 8]. The first component is a uniform distribution over $[0.78, 0.80] \times [0.2, 0.8]$. The second component has support $[0.25, 0.4] \times [0, 1]$ with X being uniform over $[0.25, 0.4]$ and Y being Beta(100, 120). The density function is therefore

$$\frac{0.35}{0.012} \times \mathbf{1}_{[0.78, 0.80] \times [0.2, 0.8]} + \frac{0.65}{0.15} \times \frac{\Gamma(220)}{\Gamma(120)\Gamma(100)} y^{99} (1-y)^{199} \mathbf{1}_{[0.25, 0.4] \times [0, 1]}. \quad (6.1)$$

Figure 6 shows the partition plots based on the exact OPT and the LL-OPT algorithms with $h = 1, \dots, 5$, respectively, for a sample size of 1000. The plots show how the root region Ω is partitioned according to the estimated tree topologies. A good partition plot should visually resemble a contour plot of the distribution to be estimated. On these partition plots, the two components of the mixture are clearly marked out by the two vertical strips located on the left-half and the right-half of the square, respectively. Furthermore, the non-uniformity of the Beta distribution in the left component is also visible through the further division of the corresponding strip.

Comparing the partition plots between the exact OPT and the LL-OPT algorithms, the partition estimated by the LL-OPT algorithm, with h as small as 2, resembles closely the partition estimated by the exact OPT algorithm with only minor differences. As h increases, the resemblance becomes even stronger, which is expected as the two algorithms should lead to exactly the same result when h is greater than the maximum depth.

The average running times used by the exact OPT algorithm and the LL-OPT algorithm with different h are summarized in Table 1, and the average Hellinger distances (computed based 100, 000 importance samples) between the estimated and the true densities are summarized in Table 2. The Dirichlet process Gaussian mixture models (DP-GMM) [Escobar and West, 1995] is also included in the comparison. All the summary statistics are based on 5 replications and corresponding standard deviations are included in the parentheses.

Based on Table 1, we find that the LL-OPT algorithm achieves substantial speedup compared with the exact OPT algorithm, especially when the sample size is large. The other algorithm for exact OPT inference, the DF-OPT algorithm, is computationally prohibitive for large sample size.

Based on Table 2, the accuracy of estimates of the LL-OPT algorithm is comparable to that of the exact OPT algorithm, with the value of h as low as 2, which is encouraging since the improvement on efficiency is sizable. For example, when $h = 2$, the LL-OPT algorithm is 15 times faster than the exact OPT algorithm for sample size of 10^5 . This makes the LL-OPT algorithm appealing for large samples. Furthermore, the estimates of the LL-OPT algorithms with $h = 2$ are almost identical, so the adaptive approach for selecting h would work well.

Example 6.2. We now consider a distribution on $[0, 1]^4$ with a bivariate Normal distribution in the first two dimensions and uniform in the other two dimensions. In this scenario, X and Y components are independent of each other and follow $\text{Norm}(0.6, 0.1^2)$ and $\text{Norm}(0.4, 0.1^2)$ respectively [Wong and Ma, 2010, adapted from Example 9], and Z and W components are independent and uniform in $[0, 1]$. As the probability mass of this distribution is essentially all located in $[0, 1] \times [0, 1]$, the density function can be approximated by

$$\frac{1}{2\pi \times 0.1^2} e^{-\frac{(x-0.6)^2 + (y-0.4)^2}{2 \times 0.1^2}}. \quad (6.2)$$

The average running times and average Hellinger distances are summarized in Tables 3 and 4 respectively. All the summary statistics are based on 5 replications and corresponding standard deviations are included in the parentheses.

Similar to the previous example, the estimation accuracy of the LL-OPT algorithm approaches that of the exact OPT algorithm as the value of h increases. Patterns of running times found in Table 3 are also similar to that of the previous example. According to Table 4, the estimation accuracy of the LL-OPT algorithm with h as low as 3 is comparable to that of the exact OPT algorithm. The improvements on running times are still sizable. When $h = 3$, the LL-OPT algorithm is over 40 times faster than the exact OPT algorithm for sample size of 10^3 , and over 190 times faster for sample size of 10^4 . Finally, estimation based on the LL-OPT algorithm stabilizes for $h = 3$.

In both simulation examples, when the sample size is large (e.g., 10^5), we can see that the running time of the LL-OPT algorithm increases roughly by a factor of p when h increases by 1, which is consistent with the time complexity analysis provided in Theorem 4.3.

6.2 Smooth density estimation

We compare the performance of the exact OPT, the LL-OPT and the FEE algorithms for density estimation to several existing methods for density estimation including kernel density estimation (KDE) [Silverman, 1986], Dirichlet process Gaussian mixture models (DP-GMM) [Escobar and West, 1995] and mixture of Pólya trees (MPT) [Lavine, 1992].

Example 6.3. We consider a strongly skewed distribution

$$x \sim \text{KT}(2, 0.1) \mathbf{1}_{[0,1]} \quad (6.3)$$

where $\Gamma(a, b) = \frac{1}{b^a \Gamma(a)} x^{a-1} e^{-\frac{x}{b}}$, and $K = 1 / \int_0^1 \Gamma(2, 0.1) dx$ is the normalizing constant.

Figure 7 shows the estimated densities by the exact OPT, the LL-OPT and the FEE algorithms for sample size $N = 10^4$. Table 5 shows the estimation errors of the three methods as well as KDE, DP-GMM and MPT measured by Hellinger distance. FEE gives the best results among all the methods.

Example 6.4. We consider a mixture of a uniform distribution and Beta distributions.

$$X_1 \sim U(0, 1); X_2 \sim \frac{4}{5}\beta(2, 10) + \frac{1}{5}\beta(7, 2) \quad (6.4)$$

Figure 8 (a) shows the densities estimated by the LL-OPT and the FEE algorithms. As a comparison, kernel density estimated is in Figure 8 (b). We see that the uniformity in the first dimension of the distribution is well captured by OPT based approaches, but not by KDE.

Example 6.5. We consider a 5-dimensional random vector

$$\mathbf{Y} = (X_1, \frac{1}{3}X_1 + \frac{2}{3}X_2, X_3, X_4, \frac{1}{5}X_3 + \frac{4}{5}X_5),$$

where $X_1, \dots, X_5$ are independent and sampled from

$$X_1 \sim \beta(2, 8), X_2 \sim \beta(8, 2), X_3 \sim U(0, 1) \quad (6.5)$$

$$X_4 \sim K_1 \Gamma(2, 1) \mathbf{1}_{[0,1]}, X_5 \sim K_2 \Gamma(1, 2) \mathbf{1}_{[0,1]} \quad (6.6)$$

where $\Gamma(a, b) = \frac{1}{b^a \Gamma(a)} x^{a-1} e^{-\frac{x}{b}}$, and $K_1 = 1 / \int_0^1 \Gamma(2, 1) dx$ and $K_2 = 1 / \int_0^1 \Gamma(1, 2) dx$ are the normalizing constants.

With a sample size of 6×10^3 , the Hellinger distances for the LL-OPT and the FEE algorithms are 0.31 and 0.27, respectively. Figure 9 shows marginal histograms based on 10^4 random samples drawn from the fitted FEE density and the true density. We can see that the fitted density resembles closely the true density in every dimension.

6.3 Real data examples

Classification is a classical application of density estimation and it covers a wide range of real world problems. We apply the LL-OPT approach to the MiniBooNE particle identification data set [Bache and Lichman, 2013] where a classifier is desired to distinguish electron neutrinos from muon neutrinos. There are a total of 130,065 data points (36,499 electron neutrinos and 93,565 muon neutrinos) in the data set, and each data point has a

dimension of 50. We divide the dataset into two equal parts, use one part for training and the other part for testing. Using the training data set, we apply LL-OPT to estimate the density for each class C_i , $i = 1, 2$ and then for each test cases x we classify it using the Bayes rule: $\hat{C}(x) = \arg \max_j P(x|C_j)P(C_j)$. We run LL-OPT with $h = 2$, and stops at regions having less than 100 data points. LL-OPT achieves a correct classification rate of 0.841, which is very close to the correct classification rate of 0.845 achieved by a support vector machine (SVM) classifier (see Section 7 for details). After applying a copula transform to make all the marginal distribution uniform, the LL-OPT achieves a even higher correct classification rate of 0.858.

7 Discussion

We provide two software packages implementing the algorithms described in this paper, available at http://www.stanford.edu/group/wonglab/opt_comp/. The **fast-opt** package implements both the exact OPT and the LL-OPT algorithms described in Section 2 and 3. The **smooth-opt** package implements the smoothing algorithm described in Section 5. Both packages are implemented in C++.

The **smooth-opt** package relies on the user to specify the value for the tuning parameter λ . An alternative approach is to use cross-validation to choose λ which requires much more computation. In our simulations we found that the result is not sensitive to the value of λ and we use $\lambda = 10^{-3}$ or 10^{-4} in all our simulations.

In our experiments, KDE was based on R package **ks** (function **kde**), DP-GMM was based on R package **bayesm** (function **rDPGibbs**), MPT was based on R package **DPpackage** (function **PTdensity**) and SVM was based on R package **e1071** (function **svm**). For all parameters we used either default values whenever available or the values suggested by the user manuals.

High dimensional density estimation is a challenging problem. The OPT is an adaptive non-parametric density estimation approach which fits to the global landscape of the data adaptively rather than locally as in kernel density estimation [Silverman, 1986], as demonstrated in our simulations. The simulations also suggest that our continuous piecewise linear FEE further improves upon the piecewise constant counterpart. The proposed FEE is a penalized regression approach. Based on a fast quadratic programming solver, it achieves decent efficiency when the sample size is reasonably high (10^5 data points). The resulting estimated density function is continuous piecewise linear, which is mostly sufficient when we want to explore how the probability mass is distributed or visualize the estimated density. When necessary, the FEE can be extended to use higher order basis functions, at the cost of heavier computation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NSF grants DMS-0906044, DMS-1407557, DMS-1330132 and NIH grant R01-HG006018. We thank IBM® for providing the academic license of Cplex®.

References

- Allaire, G. Numerical Analysis and Optimization. Oxford Science Publications; 2007.
- Bache K, Lichman M. Uci machine learning repository. 2013; 19 URL <http://archive.ics.uci.edu/ml>.
- Barron AR, Györfi L, van der Meulen EC. Distribution estimation consistent in total variation and in two types of information divergence. *IEEE Trans Inf Theor*. 1992; 38(5):1437–1454.
- Denison DG, Mallick BK, Smith AF. A bayesian cart algorithm. *Biometrika*. 1998; 85(2):363–377.
- Escobar MD, West M. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*. 1995; 90(430):577–588.
- Ferguson TS. A Bayesian Analysis of Some Nonparametric Problems. *The Annals of Statistics*. 1973; 1(2):209–230.
- Ferguson TS. Prior Distributions on Spaces of Probability Measures. *The Annals of Statistics*. 1974; 2(4):615–629.
- Freedman DA. On the asymptotic behavior of bayes' estimates in the discrete case. *The Annals of Mathematical Statistics*. 1963:1386–1403.
- Ghosh, JK.; Ramamoorthi, R. Bayesian nonparametrics. Vol. 1. Springer; 2003.
- Hutter M. Exact non-parametric bayesian inference on infinite trees. arXiv preprint arXiv:0903.5342. 2009
- Kraft CH. A class of distribution function processes which have derivatives. *Journal of Applied Probability*. 1964; 1(2):385–388.
- Lavine M. Some aspects of Polya tree distributions for statistical modelling. *The Annals of Statistics*. 1992:1222–1235.
- Lavine M. More aspects of Polya tree distributions for statistical modelling. *The Annals of Statistics*. 1994:1161–1176.
- Lo AY. On a Class of Bayesian Nonparametric Estimates: I. Density Estimates. *The Annals of Statistics*. 1984; 12(1):351–357.
- Ma L, Wong WH. Coupling optional pólya trees and the two sample problem. *Journal of the American Statistical Association*. 2011; 106(496):1553–1565.
- Mauldin RD, Sudderth WD, Williams S. Polya trees and random distributions. *The Annals of Statistics*. 1992:1203–1221.
- Parzen E. On Estimation of a Probability Density Function and Mode. *The Annals of Mathematical Statistics*. 1962; 33(3):1065–1076.
- Rosenblatt M. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*. 1956; 27(3):832–837.
- Scott DW. On optimal and data-based histograms. *Biometrika*. 1979; 66(3):605–610.
- Silverman, B. Monographs on statistics and applied probability. Chapman and Hall; 1986. Density estimation for statistics and data analysis.
- Wong WH, Ma L. Optional Pólya tree and Bayesian inference. *Ann Stat*. 2010; 38(3):1433–1459.

Figure 1.

The running time for exact OPT inference using the cached implementation (See Section 4). Running times (in seconds) are plotted against different dimensions with sample size fixed at 1,000 (left figure) as well as different sample sizes with number of dimensions fixed at 4 (right figure). The samples are distributed as bivariate normal in the first two dimensions and uniform in other dimensions. This setting helps concentrate the partitions in the first two dimensions and allow manual visual checking of the partitions (e.g., as in Figure 6).

Figure 2.
Examples of undesired situations for a triangulation.

Figure 3.
A OPT partition in 2-D (left) and its triangulation (right).

Figure 4.
A continuous piecewise linear basis function.

Figure 5.
The bottom face (the simplex Δ_j) and top face in a density plot.

Figure 6.

The partition plots based on the exact OPT and the LL-OPT algorithms with $h = 1, \dots, 5$ for Example 6.1. The samples are drawn from a mixture of uniform and “semi-beta” distributions as defined by (6.1). The sample size is 1, 000.

Figure 7.
Plot of densities estimated with 10^4 samples simulated from (6.3) by OPT and LL-OPT in red (they are complete overlapped), and FEE in black. The true density is in blue.

Figure 8.

The estimated densities for Example 6.4. (a) Densities estimated by LL-OPT (left) and FEE (right) with 10^3 and 10^4 samples simulated from (6.4) respectively. (b) Density estimated by kernel density estimation with 10^4 samples.

Figure 9.

The marginal histograms of 10^4 random samples drawn from the fitted FEE density and the true density for Example 6.5. From top to bottom are the histograms of Y_1 to Y_5 .

Table 1

Running times (in seconds) for Example 6.1. Average of 5 replicates, standard deviation reported in parentheses. The DF-OPT and DP-GMM algorithms could not finish some of the large-sample simulations within reasonable time.

Sample Size	OPT	DF-OPT	LL-OPT					DP-GMM
			$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	
10^2	0.330 (0.004)	5.856 (5.114)	0.383 (0.619)	0.386 (0.621)	0.390 (0.624)	0.394 (0.628)	0.402 (0.634)	30.120 (4.991)
	0.511 (0.008)	1229.572 (876.831)	0.404 (0.635)	0.418 (0.647)	0.434 (0.659)	0.472 (0.687)	0.559 (0.748)	244.560 (3.164)
10^3	3.442 (0.039)	* *	0.499 (0.707)	0.621 (0.788)	0.789 (0.889)	1.199 (1.095)	1.948 (1.396)	11225.040 (6.244)
	47.955 (0.121)	* *	1.885 (1.373)	3.133 (1.770)	4.946 (2.224)	9.149 (3.025)	16.354 (4.044)	* *

Table 2

Estimation errors (in Hellinger distance) for Example 6.1. Average of 5 replicates, standard deviation reported in parentheses. The DF-OPT and DP-GMM algorithms could not finish some of the large-sample simulations within reasonable time.

Sample Size	OPT	LL-OPT					DP-GMM
		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	
10^2	0.381 (0.048)	0.451 (0.056)	0.384 (0.041)	0.392 (0.045)	0.380 (0.046)	0.388 (0.050)	0.527 (0.032)
	0.183 (0.012)	0.205 (0.013)	0.183 (0.009)	0.185 (0.012)	0.180 (0.012)	0.183 (0.012)	0.424 (0.003)
10^3	0.081 (0.004)	*	0.102 (0.008)	0.082 (0.004)	0.082 (0.004)	0.081 (0.003)	0.317 (0.001)
	0.034 (0.001)	*	0.045 (0.003)	0.035 (0.002)	0.034 (0.001)	0.034 (0.001)	* *

Running times (in seconds) for Example 6.2. Average of 5 replicates, standard deviation reported in parentheses. The exact OPT algorithm could not finish the simulation with 10^5 samples in under 128GB of memory.

Table 3

Sample Size	OPT	LL-OPT					
		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
10^2	1.411	0.395	0.388	0.426	0.591	1.076	2.960
	(0.074)	(0.033)	(0.002)	(0.010)	(0.010)	(0.028)	(0.061)
10^3	31.160	0.405	0.535	0.768	1.911	6.490	25.609
	(0.532)	(0.018)	(0.048)	(0.035)	(0.035)	(0.129)	(0.432)
10^4	778.208	0.561	1.472	4.072	14.475	58.316	248.463
	(3.028)	(0.022)	(0.099)	(0.199)	(0.197)	(0.589)	(2.529)
10^5	*	2.904	12.654	40.836	154.157	615.615	2638.182
	*	(0.055)	(0.585)	(0.301)	(1.404)	(0.995)	(52.133)

Table 4

Estimation errors (in Hellinger distance) for Example 6.2. Average of 5 replicates, standard deviation reported in parentheses. The exact OPT algorithm could not finish the simulation with 10^5 samples in under 128GB of memory.

Sample Size	OPT	LL-OPT					
		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
10^2	0.561	0.922	0.612	0.571	0.564	0.562	0.561
	(0.019)	(0.005)	(0.029)	(0.029)	(0.021)	(0.022)	(0.018)
10^3	0.383	0.917	0.391	0.380	0.383	0.382	0.382
	(0.004)	(0.005)	(0.005)	(0.003)	(0.004)	(0.004)	(0.004)
10^4	0.258	0.914	0.253	0.256	0.258	0.258	0.258
	(0.001)	(0.005)	(0.001)	(0.002)	(0.003)	(0.002)	(0.001)
10^5	*	0.916	0.168	0.171	0.173	0.173	0.173
	*	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)

Table 5

Estimation errors (in Hellinger distance) of OPT, LL-OPT, FEE, KDE, DP-GMM and MPT for density (6.3). Hellinger distances estimated by averaging 10 replicates using 200, 000 importance samples for all cases. standard deviation reported in parentheses. LL-OPT runs with $h = 3$ and stops at regions with single data point. For details of other methods, see Section 7.

Sample Size	OPT	LL-OPT	FEE	KDE	DP-GMM	MPT
10^2	0.1793	0.1808	0.1421	0.1502	0.1716	0.1662
	(0.0377)	(0.0413)	(0.0274)	(0.0147)	(0.0150)	(0.0190)
10^3	0.0813	0.0813	0.0476	0.0861	0.0868	0.0760
	(0.0098)	(0.0098)	(0.0077)	(0.0058)	(0.0094)	(0.0073)
10^4	0.0346	0.0340	0.0187	0.0483	0.0569	0.0440
	(0.0035)	(0.0029)	(0.0030)	(0.0025)	(0.0059)	(0.0058)