

Hidden Markov Pólya trees for high-dimensional distributions

Naoki Awaya and Li Ma

Department of Statistical Science

Duke University, Durham, NC 27708, USA

December 9, 2021

Abstract

The Pólya tree (PT) process is a general-purpose Bayesian nonparametric model that has found wide application in a range of inference problems. It has a simple analytic form and the posterior computation boils down to beta-binomial conjugate updates along a partition tree over the sample space. Recent development in PT models shows that performance of these models can be substantially improved by (i) allowing the partition tree to adapt to the structure of the underlying distributions and (ii) incorporating latent state variables that characterize local features of the underlying distributions. However, important limitations of the PT remain, including (i) the sensitivity in the posterior inference with respect to the choice of the partition tree, and (ii) the lack of scalability with respect to dimensionality of the sample space. We consider a modeling strategy for PT models that incorporates a flexible prior on the partition tree along with latent states with Markov dependency. We introduce a hybrid algorithm combining sequential Monte Carlo (SMC) and recursive message passing for posterior sampling that can scale up to 100 dimensions. While our description of the algorithm assumes a single computer environment, it has the potential to be implemented on distributed systems to further enhance the scalability. Moreover, we investigate the large sample properties of the tree structures and latent states under the posterior model. We carry out extensive numerical experiments in density estimation and two-group comparison, which show that flexible partitioning can substantially improve the performance of PT models in both inference tasks. We demonstrate an application to a mass cytometry data set with 19 dimensions and over 200,000 observations.

1 Introduction

The Pólya tree (PT) (Freedman, 1963; Ferguson, 1974; Lavine, 1992) is a stochastic process that generates random probability measures and is introduced as a prior for Bayesian nonparametric inference. While the PT generalizes the Dirichlet process (DP) (Ferguson, 1973) as it yields the DP under certain hyperparameters (Ferguson, 1974), the statistical properties and practical applications of the PT are very different. While the DP is most frequently used as a mixing distribution that induces clustering structures, the PT is often adopted for directly modeling probability densities.

The PT is defined generatively on a recursive partition—or a partition tree—over the sample space through coarse-to-fine sequential probability assignment at each tree split. In a classical (univariate) PT, the tree is dyadic and the conditional probability assigned to the two children nodes at each tree split arises from independent beta priors, which leads to analytic simplicity and ease in computing the posterior. Obtaining the posterior is straightforward from beta-binomial conjugacy and incurs a computational budget that scales only linearly with the sample size, making the PT one of the few nonparametric models applicable to data with massive sample size. Moreover, the posterior computation is embarrassingly parallelizable across the tree nodes.

The PT has been applied in various contexts beyond the original application of density estimation. A far-from-exhaustive list includes survival analysis (Muliere and Walker, 1997; Neath, 2003), imputing missing values (Paddock, 2002), goodness-of-fit tests (Berger and Guglielmi, 2001), two-group comparison (Ma and Wong, 2011; Holmes et al., 2015; Chen and Hanson, 2014; Soriano and Ma, 2017), density regression (Jara and Hanson, 2011), ANOVA (Ma and Soriano, 2018), testing independence (Filippi et al., 2017), and hierarchical modeling (Christensen and Ma, 2020). The PT has also been utilized in semi-parametric analyses such as in (generalized) linear models (Walker et al., 1999; Hanson and Johnson, 2002; Walker and Mallick, 1997).

Early developments of the PT are based on an *a priori* fixed partition tree on the sample space. The resulting inference can be sensitive to the choice of the partition points defining the tree. In particular, the resulting process, both *a priori* and *a posteriori*, can be jumpy at these points. In hypothesis testing and model choice, this sensitivity is also reflected in the sometimes substantial change in the marginal likelihood/Bayes factor when the partition points are slightly varied. To remedy the issue, Paddock et al. (2003) modified the PT model so that observations are generated from the PT model with slightly different partition points. Hanson and Johnson (2002) and Hanson (2006) proposed a mixture of PTs by defining partition points along quantiles of a parametric model endowed with a hyperprior to allow model averaging on the partition points. This strategy does not allow individual partition points to adapt to local features of the distribution but only the whole set of points to the global structure of the distribution, and is most effective when the underlying density is close to the specified parametric model. Nieto-Barajas and Müller (2012) took a different approach by modeling the probability assignments within each level of the tree in a correlated manner to smooth out the random measure over the boundaries of partitioning. While these approaches alleviate the sensitivity to partition points in low-dimensional settings, they are not easily applicable (though in principle possible) to problems with even just a handful of dimensions. Moreover, Bayesian inference with these models generally require MCMC, whose effectiveness can (in fact often does) still suffer from the sensitivity with respect to the partition points.

Another related issue regarding the partitioning scheme of the PT is that in multivariate problems, traditionally the partition tree is constructed by dividing all dimensions of the sample space at each split. For example, for a d -dimensional sample space, each time a tree node is divided, it is split into 2^d children nodes, and probability assignment over these 2^d child nodes is modeled by independent Dirichlet priors.

Wong and Ma (2010) noted that such a “symmetric” partition scheme is undesirable as the dimensionality increases, in which case due to the exponential growth of the partition blocks, the vast majority of the blocks are barely, if at all, populated by data. As such they propose to incorporate adaptivity into the partitioning strategy with respect to the structure of the underlying distribution through adopting a Bayesian CART-like prior (Chipman et al., 1998) on the space of dyadic partition trees.

However, in order to maintain the analytic simplicity of the posterior and achieving MCMC-free exact Bayesian inference with a linear computational budget, the Bayesian CART-like prior has to be restricted to only divide at the middle point (or otherwise a pre-determined fixed point) on one of the dimensions at each tree split. Not only does this hamper the model’s ability to adapt to distributional structures, but it makes the model suffer from the same sensitivity with respect to the partition points. Moreover, even with this restriction, the inference algorithm (based on recursive message passing) is only computationally practical for up to about 10 dimensions on continuous sample spaces.

In a different vein, recent developments have demonstrated that aside from enhancing the partitioning strategy, the PT can also be substantially improved by adopting more flexible priors (as opposed to independent betas) on the probability assignment at each tree split (Jara and Hanson, 2011; Nieto-Barajas and Müller, 2012; Ma, 2017). One strategy for enriching the PT in this regard is by introducing latent state variables at each tree split and adopt priors on the probability assignment *given* these states. When the latent states are discrete and modeled with Markov dependency, analytical simplicity is preserved and exact Bayesian inference can proceed through recursive message passing that maintains the linear computational budget (Ma, 2017).

Given these developments, we are motivated by the following questions: Is it possible to incorporate into the PT a very flexible partition tree prior, such as the general

Bayesian CART (i.e., without the restriction to partition at middle points), that will (i) enhance its adaptivity to distributional structures in multivariate settings; (ii) resolve its sensitivity to the choice of partition points; and (iii) allow a tractable form of the joint posterior and a posterior inference algorithm that is scalable to moderately high-dimensional problems (e.g., up to 100 dimensions)? Moreover, should such a strategy exist, can the resulting model and inference algorithm be made compatible with incorporating (possibly Markov dependent) latent states on the tree nodes?

The goals of making the partition tree prior more flexible while enhancing the computational scalability appear at odds with each other. Large tree spaces are well known to be very hard to compute over. In moderate to high dimensional settings exact inference involving flexible tree structures is beyond reach and even the most advanced MCMC approaches tailor-made for trees encounter substantial difficulty due to the pervasive multi-modality of distributions in such spaces. Recent advances in sequential Monte Carlo (SMC) for regression tree models (Lakshminarayanan et al., 2013; Lu et al., 2013), however, suggest that efficient inference is possible in moderately high-dimensional settings (up to about 100 dimensions). Moreover, once the partition tree is sampled, the conditional posterior for the rest of the model can be computed analytically through recursive message passing. We will therefore exploit a hybrid strategy that uses a new SMC sampler to efficiently sample from the marginal posterior of the partition tree structure, along with recursive message passing to compute the exact conditional posterior of the latent state variables given the tree.

Beyond the methodological development, we will also investigate the theoretical properties of the posterior on the partition tree and the latent states. Previous theoretical literature on the PT and related models have mostly focused on establishing the posterior consistency and the contraction rate of the random measure induced under these models (Castillo, 2017a; Castillo and Randrianarisoa, 2021). In multivariate

settings, however, the partition tree itself is highly informative about the underlying distribution. Moreover, in applications involving model choice and hypothesis testing, it is often the latent states, not the random measures, that are of direct interest. As such, we focus on studying the asymptotic behavior of the marginal posterior on the partition tree and latent states, establishing consistency results on their convergence toward the trees and states that most closely characterize the underlying truth.

The rest of the paper is organized as follows. In Section 2 we describe a flexible prior on the partition tree structure that relaxes the restriction of “dividing in the middle” on partition points and present a general form of PT models that adopt this prior along with latent states associated with the tree nodes with a Markov dependency structure. In Section 3, we present our hybrid computational strategy that can work effectively up to 100 dimensions consisting of an SMC algorithm for sampling on the marginal posterior of the partition tree and a recursive message passing algorithm for obtaining the exact conditional posterior of the latent states and the random measure given the sampled trees. In Section 4 we investigate the asymptotic properties of the tree structures and latent states identified under the posterior model. In Section 5, we carry out extensive numerical experiments to examine the performance of our method in the context of two important applications of PTs—density estimation and the two-group comparison, followed by an application to a data set from mass cytometry in Section 6. In Section 7 we conclude with a brief discussion. All proofs are provided in **Supplementary Materials C**.

2 Method

We first review the PT model (Ferguson, 1973; Lavine, 1992) on a dyadic recursive partition in Section 2.1. The model, while defined on a general multivariate sample space, differs from a traditional multivariate PT which adopts a multi-way symmetric

recursive partitioning. Then we introduce a new class of PT models that incorporates both the flexible partition prior and latent states with Markov dependency.

2.1 Pólya trees defined on recursive dyadic partitions

Without loss of generality, we consider a continuous sample space represented as a d -dimensional rectangle $\Omega = (0, 1]^d$. For unbounded sample spaces such as \mathbb{R}^d , one can transform each margin to $[0, 1]$ by applying, say, a cumulative distribution function transform or by standardizing the data based on its observed range of values. We use μ to denote the Lebesgue measure on Ω . A (dyadic) recursive partitioning T on Ω is a sequence of partitions of Ω such that the partition blocks at each level of the partitioning are obtained by dividing each block in the previous level into two children blocks. Formally, we can write $T = \bigcup_{k=0}^{\infty} \mathcal{A}^k$, where \mathcal{A}^k is a partition of Ω in the k th level. More specifically, $\mathcal{A}^0 = \{\Omega\}$, and $A \in \mathcal{A}^k$ ($k = 0, 1, 2, \dots$) is divided into A_l and A_r , which satisfy $A_l, A_r \in \mathcal{A}^{k+1}$, $A_l \cup A_r = A$, and $A_l \cap A_r = \emptyset$. (Throughout the paper, a subscript “ l ” or “ r ” on a node indicates the left or right child node.) For example, when $d = 1$ and if the tree is recursively divided at the middle point of each node, then nodes in level k are of the form $(l/2^k, (l+1)/2^k]$ for some $l \in \{0, \dots, 2^k - 1\}$. Another common strategy is to define the tree based on the quantiles of a probability measure F so that $A \in \mathcal{A}^k$ is of the form $A = (F^{-1}(\frac{l}{2^k}), F^{-1}(\frac{l+1}{2^k})]$ for $l \in 0, \dots, 2^k - 1$.

Given a partition tree T , we can define a random measure Q by putting a prior on the conditional probability $\theta(A) = Q(A_l|A) = 1 - Q(A_r|A)$ at each $A \in T$. Under the PT prior, the parameters $\theta(A)$ follow independent beta distributions $\text{Beta}(\alpha_l(A), \alpha_r(A))$, where $\alpha_l(A)$ and $\alpha_r(A)$ are hyperparameters. The corresponding posterior, given an i.i.d. sample x_1, \dots, x_n from Q , is again a PT with a simple conjugate update on the conditional probabilities:

$$\theta(A) \mid x_1, \dots, x_n \sim \text{Beta}(\alpha_l(A) + n(A_l), \alpha_r(A) + n(A_r)),$$

where $n(A)$ represents the number of observations in a set $A \subset \Omega$. Though the tree needs to be infinitely deep to ensure full support of the PT, for practical purposes, one typically sets a sufficiently large maximum depth (or resolution) of T and compute the posteriors of $\theta(A)$'s defined on this finite tree structure (Hanson and Johnson, 2002). We shall refer to a node in the deepest level as a “leaf” or “terminal node”. On a leaf, the conditional distribution can be set to a baseline $F(\cdot|A)$, such as the uniform distribution $\mu(\cdot|A)$. In Section 3 when we present inference algorithms, we shall adopt this practical strategy and assume T is finite and use $\mathcal{N}(T)$ and $\mathcal{L}(T)$ to denote the collection of the non-terminal nodes and the leaf nodes, respectively.

2.2 Incorporating flexible partition points

We incorporate a Bayesian-CART like prior on T by randomizing both the dimension in which to divide a node and the location to divide. Our prior relaxes the “always-divide-in-the-middle” restriction imposed in Wong and Ma (2010). This prior on the partition tree T differs from that in the mixture of PTs of Hanson (2006), which does not randomize over the dimension to divide, but generates the boundaries of the tree nodes jointly using quantiles of a parametric family.

Our prior can be described iteratively as a generative process that recursively divides the sample space. Specifically, suppose we have a node A in the rectangular form, $A = (a_1, b_1] \times \cdots \times (a_d, b_d]$. We divide A into two rectangular children by cutting along a randomly chosen dimension at a random location. The dimension to divide $D(A) \in \{1, 2, \dots, d\}$, and the (relative) location to divide $L(A) \in (0, 1)$ are given independent priors of the following forms:

$$D(A) \sim \text{Mult}(\lambda_1(A), \dots, \lambda_d(A)) \quad \text{and} \quad L(A) \sim \sum_{l=1}^{N_L-1} \beta_l(A) \delta_{l/N_L}(\cdot), \quad (1)$$

where $\delta_x(\cdot)$ represents the unit point mass at x , and $N_L - 1$ is the total number of grid points along $(0, 1)$. Both $\{\lambda_i(A)\}_{i=1, \dots, d}$ and $\{\beta_l(A)\}_{l=1, \dots, N_L-1}$ sum to 1. In the above, we have adopted a uniform grid over $(0, 1)$ for notational simplicity, but it

does not have to be as such. With $D(A) = j$ and $L(A) = l/N_L$, the two children nodes A_l and A_r are

$$A_l = (a_1, b_1] \times \cdots \times (a_j, a_j + l/N_L \cdot (b_j - a_j)] \times \cdots \times (a_d, b_d],$$

$$A_r = (a_1, b_1] \times \cdots \times (a_j + l/N_L \cdot (b_j - a_j), b_j] \times \cdots \times (a_d, b_d].$$

In principle one could adopt a continuous prior on the partition location $L(A)$. A discretized prior is helpful, however, because it will substantially simplify posterior computation. In practice, as long as the grid is dense enough, the discrete prior will be practically just as flexible. Indeed we have verified in extensive numerical experiments that when N_L is large enough (more than 30 to 50) over a uniform grid, posterior inference no longer improves in any noticeable way.

For the prior on $D(A)$, we set $\lambda_j(A) = 1/d$ for all nodes A as a default choice. When $L(A)$ is given a weak prior widely spread over $(0, 1)$, the resulting inference can be sensitive to the “tail” behavior of the distribution in the node, resulting in high posteriors of $L(A)$ near the extreme values 0 and 1. A detailed discussion on this phenomenon will be provided in Section 5.1.1. This issue can be effectively addressed by making the prior of $L(A)$ depend on the sample size $n(A)$ so that it encourages more balanced divisions at large sample sizes. More specifically, we adopt the following prior with an exponentially decaying tail

$$P(L(A) = l/N_L) = \beta_l \propto \exp[-\eta n(A) f(|l/N_L - 0.5|)], \quad l = 1, \dots, N_L - 1, \quad (2)$$

where $\eta \geq 0$ is a hyperparameter and $f : [0, \infty) \rightarrow [0, \infty)$ is an increasing function with $f(0) = 0$. In the following, we shall use a function $f(x) = x$, and so our prior on $L(A)$ is a (discretized) Laplace distribution. We provide theoretical justification for adopting this prior with exponential tails in Section 4.

Another generalization of the prior on $L(A)$ is to incorporate a spike-and-slab set-up with a spike at the middle point $1/2$. In particular, one can adopt a dependent spike prior among the nodes such that once a node A is divided exactly at the

middle point, so are its descendants. This generalization will substantially reduce the amount of computation in regions of the sample space where the data are either sparse or lack interesting structure, e.g., close to the uniform distribution. We implement the spike-and-slab in our software but defer the details of this generalization to **Supplementary Materials A** to avoid distracting the reader from the main ideas.

Given the tree prior, our PT model now consists of the two components—tree generation and conditional probability assignment. Figures 1(a) and 1(b) present a graphical model representation for each.




Figure 1: Graphical representation of PT models given the tree T . The hyperparameters are hidden for simplicity.

2.3 Hidden Markov Pólya tree models

2.3.1 General framework

Next we extend the above model to accommodate two recent developments in the PT literature: (i) incorporating latent state variables along the tree structure and (ii) joint modeling of multiple groups of observations. Incorporating latent variables allow more flexibly characterizing distributional features through adding prior dependency. As in recent literature, we consider incorporating discrete state variables that follow a Markov process along the tree structure. Because the description in this section always pertains to the model *given* the randomly generated partition tree T , for brevity we shall not keep stating “given T ”.

We generalize our notation to allow observing one or more groups of i.i.d. observations. Let G be the number of groups of i.i.d. observations. For the g th group

($g = 1, 2, \dots, G$), let Q_g be the sampling measure for that group. Let \mathbf{Q} denote the collection of all G sampling measures. That is, $\mathbf{Q} = \{Q_g\}_{g=1}^G$. Let $\mathbf{x}_g = (x_{g,1}, \dots, x_{g,n_g})$ denote the observations in the g th group, which are i.i.d. given Q_g , where n_g , the sample size for the group, is allowed to differ across the groups. We use $\mathbf{x} = \{\mathbf{x}_g\}_{g=1}^G$ to denote the collection of all observations from all groups.

Next we specify a prior on \mathbf{Q} in terms of a joint prior on the conditional probability on each $A \in T$, $\theta_g(A) = Q_g(A_l | A) = 1 - Q_g(A_r | A)$. We use latent variable modeling to incorporate prior dependency among the tree nodes. Specifically, let $\{V(A) : A \in T\}$ denote a collection of latent state variables, one for each A , and without loss of generality, assume that $V(A)$ takes discrete values from $\{1, \dots, I\}$. (In practice, the number of states I can differ among A .) Joint priors of $\theta_g(A)$ for all g and A are then defined conditionally on these latent states.

Existing literature has exploited these latent states to characterize both the within-group structure of each distribution Q_g and the between-group relationship among the Q_g . An example of within-sample structures is the smoothness of each underlying distribution, which is explored in the context of density estimation (Ma, 2017). An example of between-group structures is the difference between two (or more) distributions (Soriano and Ma, 2017).

Dependent modeling of the latent states over the partition tree is desirable as *a priori* one would expect interesting structures (both within-group and between-group) to exhibit themselves in a correlated manner over the sample space. For example, functions tend to have similar smoothness over adjacent locations, and two-group difference tend to be clustered in space. A computationally efficient strategy for modeling such dependency over the tree is by a hidden Markov process along the tree (Crouse and Baraniuk, 1997), which starts from the root node, $A = \Omega$, and sequentially generates the latent states in a coarse-to-fine fashion according to

(possibly node-specific) transition matrices $\boldsymbol{\xi}(A)$ whose (i, i') th element is

$$\boldsymbol{\xi}_{i,i'}(A) = P(V(A) = i' \mid V(A^p) = i) \quad \text{for } i, i' \in \{1, \dots, I\},$$

where A^p denotes A 's parent. (We shall use superscript “ p ” to indicate the parent of a node in T .) For $A = \Omega$, since Ω has no parent, we can simply let $\boldsymbol{\xi}_{i,i'}(\Omega)$ be constant over i , representing the initial state probabilities on Ω .

Given the $V(A)$'s, $\{\theta_g(A)\}_{g=1,\dots,G}$ can then be modeled as conditionally independent *a priori*. The specific choices of these conditional priors are problem-dependent. We give two examples below. Figure 1(c) presents a graphical model representation for the latent state modeling on G probability distributions by PTs given T , which along with our generalized prior on the partition tree T presented in Figure 1(a) forms the most general version of the model we consider in this work.

Example 1: Density estimation with adaptive smoothness

An example of within-group structures that the latent state $V(A)$ can characterize is the smoothness of the density functions for the random measures. For example, Ma (2017) proposed the adaptive Pólya tree (APT) model which incorporates latent states to allow different levels of local smoothness in the underlying distribution. This is achieved by modeling the $\theta(A)$'s as $\text{Beta}(m(A)\nu(A), (1 - m(A))\nu(A))$, where $m(A)$ is the prior mean and $\nu(A)$ the precision parameter which characterizes the smoothness of the random measure with larger $\nu(A)$ corresponding to more smoothness, and then modeling the precision parameters conditional on the latent state $V(A)$ with a hyperprior $\nu(A) \mid V(A) = i \sim F_i$ with the (conditional) hyperprior F_i given $V(A) = i$. The hyperpriors F_1, F_2, \dots, F_I are chosen to be stochastically increasing $F_1 \prec F_2 \prec \dots \prec F_I$. That is, $F_i([x, \infty)) \leq F_{i+1}([x, \infty))$ for all $x > 0$. The rate at which the latent states transition into a higher latent state along each subbranch of the partition tree determines the local smoothness of the underlying density.

Example 2: Two-group comparison

In two-group comparison, we are interested in testing and identifying differences between two measures $\mathbf{Q} = \{Q_g\}_{g=1,2}$ based on an i.i.d. sample from each. The “global” testing problem can be formulated as testing the following null and alternative hypotheses: $H_0 : Q_1 = Q_2$ vs $H_1 : Q_1 \neq Q_2$. Noting that two-group differences may exist in parts of the sample space and not others, the coupling OPT (Ma and Wong, 2011) and the multi-resolution scanning (MRS) model (Soriano and Ma, 2017) are PT-based models that allow the measures to differ on some nodes $A \in T$ and not others. This more “local” perspective on two-group comparison enables these models to not only test for H_0 vs H_1 , but to identify regions on which the two measures differ. To achieve this, these models incorporate state variables that characterize whether the conditional probabilities on each A are equal:

$$V(A) = 1 \Leftrightarrow Q_1(A_l | A) \neq Q_2(A_l | A), \quad V(A) = 2 \Leftrightarrow Q_1(A_l | A) = Q_2(A_l | A). \quad (3)$$

When $V(A) = 1$, the two corresponding conditional probabilities are given independent beta priors, whereas if $V(A) = 2$, they are tied and given a single beta prior. Markov dependency among the states on different nodes are incorporated to induce the desired spatial correlation of cross-group differences. Additional latent states can be further incorporated to reflect more complex relationships between the distributions. In fact, the MRS also incorporates an additional state $V(A) = 3$, which introduces the same coupled prior as $V(A) = 2$, but works as an absorbing state that once $V(A) = 3$, all descendants of A will remain in that state, corresponding to the case that the conditional distributions $Q_1(\cdot|A)$, and $Q_2(\cdot|A)$ are completely equal.

3 Bayesian inference

In sum, the models we consider all share a common structure consisting of the following components: (i) the partition tree T defined by the dimension and location

variables D 's and L 's, which follow the priors given in Eq. (1); (ii) the latent state variables $V(A)$ given T which follow a Markov prior; (iii) the conditional probabilities along the given tree T , $\{\theta_g(A)\}_{g=1}^G$, whose joint prior are specified independently across the nodes on T given the latent states; and finally (iv) given the random measures Q_g defined by T and $\theta_g(A)$'s, we observe an i.i.d. sample \mathbf{x}_g from each Q_g , independently across g . Formally, we have the following full hierarchical model:

$$\begin{aligned}
T &| \boldsymbol{\lambda}, \eta \sim p(T | \boldsymbol{\lambda}, \eta) \\
\{V(A) : A \in T\} &| \boldsymbol{\xi}, T \sim \text{Markov}(\boldsymbol{\xi}) \\
(\theta_1(A), \dots, \theta_G(A)) &| V(A), T \stackrel{\text{ind}}{\sim} p(\theta_1(A), \dots, \theta_G(A) | V(A)) \text{ for } A \in T \\
\mathbf{x}_g = (x_{g,1}, x_{g,2}, \dots, x_{g,n_g}) &| Q_g \stackrel{\text{iid}}{\sim} Q_g \text{ for } g = 1, 2, \dots, G.
\end{aligned}$$

The key to Bayesian inference is the ability to either compute or sample from the joint posterior $(T, \mathbf{V}, \boldsymbol{\theta})$ given all data $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_G)$, where \mathbf{V} and $\boldsymbol{\theta}$ represent the totality of all latent states and conditional probabilities given T respectively. While in some problems such as density estimation one may mainly be interested in just the marginal posterior of the Q_g 's, in others such as two-group comparison where one wants to characterize the between-group relationships among the distributions, the latent states (along with T), which characterizes such relationships, are often of prime interest. In multivariate and even high-dimensional problems, the tree structure T is also of great interest as it sheds light on the underlying structures in the distributions.

To this end, we shall take advantage of recent developments in sequential Monte Carlo (SMC) sampling for tree-based models (Lakshminarayanan et al., 2013; Lu et al., 2013) and advances in message passing algorithms for PT models with Markov dependency (Ma, 2017). We introduce a hybrid algorithm that combines these two computational strategies to effectively sample from the joint posterior in high-dimensional spaces. Overall, the hybrid algorithm consists of two stages:

1. Sampling from the marginal posterior of the partition tree

We design an SMC sampler—that is, a particle filter—to sample a collection of tree structures T^1, \dots, T^M by growing each tree from coarse to fine scales. It uses one-step look-ahead message passing to construct proposal distributions for $D(A)$ and $L(A)$, one node at a time.

2. Computing the conditional posterior given the sampled trees

Given each tree sampled by the SMC, we analytically compute the exact conditional posteriors of $V(A)$'s and $\theta(A)$'s using recursive message passing. (The detailed algorithm will be given in Section 3.2.)

3.1 SMC to sample from tree posterior

In the SMC stage to sample the trees, each particle stores a realized form of a finite tree structure, and one node of each tree is divided at each step of the SMC sampling. Suppose T_t is the finite tree obtained after dividing the sample space t times in a particle, and for this tree we define the target distribution

$$\pi_t(T_t) = P(T_t | \mathbf{x}) \propto P(T_t)P(\mathbf{x} | T_t).$$

Here $P(T_t)$ is the joint prior of the variables $D(A)$'s and $L(A)$'s for the non-leaf nodes of T_t , and $P(\mathbf{x} | T_t)$ is the marginal likelihood given the tree T_t under the hierarchical model, in which \mathbf{V} and $\boldsymbol{\theta}$ are integrated out. To sample from this target distribution, we sequentially construct a set of M particles $\{T_t^m, W_t^m\}_{m=1}^M$, where T_t^m is a realized tree and W_t^m is the associated importance weight for the m th particle. Examples of generated trees are given in Figure 2, where the sample space has been divided three times, and in the next step, new partition boundaries will be added in the gray nodes.

Following Lakshminarayanan et al. (2013), we adopt in each step of the SMC a breadth-first tree-growth strategy by dividing the “oldest active” leaf node—that is, the one generated in the earliest step and is yet to be terminated in division. Further division of a node is terminated once the number of observations in that

node is below a pre-set threshold (e.g., 5 in our software implementation) to avoid overfitting. Otherwise a node is bisected along a boundary whose dimension and location are randomly drawn from a proposal distribution. For each particle, a finite tree T_t is formed by a sequence of decisions $\{J_s\}_{s=1}^t$, where $J_s = (D_s, L_s)$ correspond to all of the variables $D(A)$ and $L(A)$ at the s th step of the SMC.




Figure 2: An example of realized finite trees in the particle system obtained after the step $t = 3$. The numbers in the squares indicate in which step the boundaries are drawn. Among the current leaf nodes, the nodes colored gray are the oldest nodes generated in the earliest step, so they are split in the next step.

As we will see in Proposition 3.1, the target distribution $\pi_t(T_t)$ has a decomposition $\pi_t(T_t) = C_t \pi_{t-1}(T_{t-1}) \pi_t(J_t | T_{t-1}) w_t(T_{t-1})$, where C_t is a constant independent of T_t , $\pi_t(J_t | T_{t-1})$ a conditional distribution on J_t given T_{t-1} , and $w_t(T_{t-1})$ a function of T_{t-1} . We will choose $\pi_t(J_t | T_{t-1})$ as the proposal for J_t under which the corresponding importance weight will simply be $w_t(T_{t-1})$, independent of J_t .

More specifically, suppose at the current step t , we are to divide $A_t \in T_{t-1}$, into $A_{t,l}$ and $A_{t,r}$ with decision J_t . Let $M_i(A_t | J_t)$ be the marginal likelihood on the node A_t under the decision J_t evaluated based on the observations in A_t . That is,

$$M_i(A_t | J_t) = \int \left[\prod_{g=1}^G \theta_g(A_t)^{n_g(A_{t,l})} (1 - \theta_g(A_t))^{n_g(A_{t,r})} \right] dP(\theta_1(A_t), \dots, \theta_G(A_t) | V(A_t) = i), \quad (4)$$

where $n_g(A)$ is the number of observations of the g th group included in A . To avoid cumbersome notation, we suppress in our notation the dependency of $M_i(A_t | J_t)$ on the observations \mathbf{x} . For example, if the $\{\theta_g(A_t)\}_{g=1}^G$ follow independent beta priors

written as $\text{Beta}(\alpha_l^i(A_t), \alpha_r^i(A_t))$ given $V(A_t) = i$, then the marginal likelihood has the following expression $M_i(A_t | J_t) = \prod_{g=1}^G \frac{B(\alpha_l^i(A) + n_g(A_{t,l}), \alpha_r^i(A) + n_g(A_{t,r}))}{B(\alpha_l^i(A_t), \alpha_r^i(A_t))}$, where $B(\cdot, \cdot)$ is the beta function. Based on the values of $M_i(A_t | J_t)$, we can analytically compute the proposal and the importance weight using a general recursive algorithm, as described in the following proposition.

Proposition 3.1. *For every possible decision J_t and states $i = 1, \dots, I$, let $\varphi_i(A_t)$ be a function defined recursively:*

$$\varphi_i(A_t) = \begin{cases} \frac{\xi_{1,i}(\Omega)M_i(\Omega|J_t)}{\sum_{j=1}^I \xi_{1,j}(\Omega)M_j(\Omega|J_t)} & \text{if } A_t = \Omega \\ \frac{\sum_{j=1}^I \varphi_j(A_t^p) \xi_{j,i}(A_t) M_i(A_t|J_t)}{\sum_{k=1}^I \sum_{j=1}^I \varphi_j(A_t^p) \xi_{j,k}(A_t) M_i(A_t|J_t)} & \text{Otherwise,} \end{cases} \quad (5)$$

where A_t^p is A_t 's parent node. Also, let $h(J_t | A_t)$ be a function of J_t defined as

$$h(J_t | A_t) = \sum_{i=1}^I \left\{ \sum_{j=1}^I \varphi_j(A_t^p) \xi_{j,i}(A_t) \right\} M_i(A_t | J_t) \frac{\mu(A_{t,l})^{-n(A_{t,l})} \mu(A_{t,r})^{-n(A_{t,r})}}{\mu(A_t)^{-n(A_t)}}, \quad (6)$$

where $n(A)$ denotes the total number of observations included in a node A . Then the target distribution $\pi_t(T_t)$ can be expressed in terms of $\pi_{t-1}(T_{t-1})$ as

$$\pi_t(T_t) = C_t \pi_{t-1}(T_{t-1}) \pi_t(J_t | T_{t-1}) w_t(T_{t-1}),$$

where C_t is a constant and

$$\pi_t(J_t | T_{t-1}) = \frac{P(J_t) h(J_t | A_t)}{\sum_{j_t} P(j_t) h(j_t | A_t)}, \quad w_t(T_{t-1}) = \sum_{j_t} P(j_t) h(j_t | A_t)$$

The summation over j_t is taken over all possible decisions.

Corollary 3.1. *Let $h(J_t | A_t)$ be the function defined in Proposition 3.1. Then the proposal distribution $\pi_t(D_t | T_{t-1})$ is given by*

$$\pi_t(D_t | T_{t-1}) = \pi_t(D_t | T_{t-1}) \pi_t(L_t | D_t, T_{t-1}), \quad \text{where}$$

1. $\pi_t(D_t | T_{t-1})$ is $\text{Mult}(\tilde{\lambda}_1(A_t), \dots, \tilde{\lambda}_d(A_t))$ with

$$\tilde{\lambda}_j(A_t) \propto \sum_{l=1}^{N_L-1} \pi_t((j, l/N_L) | T_{t-1}) \propto \lambda_j(A_t) \sum_{l=1}^{N_L-1} \beta_l(A_t) h((j, l/N_L) | A_t).$$

2. Given $D(A_t) = j$, the conditional posterior of $L(A_t)$ is

$$\pi_t(L_t = l/N_L | D_t = j, T_{t-1}) = \sum_{l=1}^{N_L-1} \tilde{\beta}_l(A_t) \delta_{l/N_L}(\cdot),$$

for $j = 1, 2, \dots, I$ and $l = 1, \dots, N_L - 1$ with

$$\tilde{\beta}_l(A_t) \propto \beta(A_t) h(j, l/N_L | T_{t-1}).$$

We also have an analytical expression of the incremental weight:

$$w_t(T_{t-1}) = \sum_{j=1}^d \sum_{l=1}^{N_L-1} \lambda_j(A_t) \beta_l(A_t) h((j, l/N_L) | A_t).$$

Remark: The recursive function $\varphi_i(A_t)$ can be computed based on $\varphi_i(A_t^p)$ with the fixed computational cost. Hence, the optimal proposal $\pi_t(J_t | T_{t-1})$ and the incremental weight $w_t(T_{t-1})$, which are functions of $h(J_t | A_t)$, can be obtained at each step with constant computational cost with complexity $O(I^2 N_L d n(A_t))$. As such, our inference algorithm scales linearly in both the dimensionality and the sample size.

We summarize the algorithm in updating the particle system from $\{T_{t-1}^m, W_{t-1}^m\}_{m=1}^M$ to $\{T_t^m, W_t^m\}_{m=1}^M$ below. All operations are repeated for $m = 1, \dots, M$.

1. Choosing the current node

From T_{t-1}^m , choose the the node generated in the earliest step among the current leaf nodes that have not been terminated, which is denoted by A_t .

2. Obtaining the information of the parent node

Locate A_t 's parent node, A_t^p , and fetch the values of $\varphi_i(A_t^p)$ for $i = 1, \dots, I$.

3. Computing the necessary quantities

For all possible $J_t = (D_t, L_t)$, compute $M_i(A | J_t)$ ($i = 1, \dots, I$) and $h(J_t | A_t)$.

4. Dividing the current node

Compute the parameters $\tilde{\lambda}_j(A_t)$ for $j = 1, \dots, d$ and sample

$$D_t^m \sim \text{Mult}(\tilde{\lambda}_1(A_t), \dots, \tilde{\lambda}_d(A_t)).$$

Given D_t^m , compute the parameters $\tilde{\beta}_l(A_t)$ for $l = 1, \dots, N_L - 1$ and sample

$$L_t^m \sim \sum_{l=1}^{N_L-1} \tilde{\beta}_l(A_t) \delta_{l/N_L}(\cdot).$$

Divide the current node A_t with $J_t^m = (D_t^m, L_t^m)$ to update the tree T_t^m .

5. Updating the importance weight

Compute the incremental weight $w_t(T_{t-1}^m)$ and update the importance weights

$$W_t^m = \frac{W_{t-1}^m w_t(T_{t-1}^m)}{\sum_{m'=1}^M W_{t-1}^{m'} w_t(T_{t-1}^{m'})}.$$

If the effective sample size $1 / \sum_{m=1}^M (W_t^m)^2$ is less than some prespecified threshold (e.g., $M/10$), resample the particles.

6. Computing the information on the current node for its descendants

Given J_t^m , compute $\varphi_i(A_t)$ for $i = 1, \dots, I$.

In the algorithm, we stop dividing A_t if either the depth of A_t is equal to a pre-set maximum resolution K (e.g., 15) or the number of observations in A_t is less than a pre-set threshold (e.g., 5). The SMC algorithm terminates when all the nodes of all the particles have been stopped. The maximum resolution K controls the level of local details that the model allows to infer, and larger values of K require more computational time. In a wide range of applications we have found that setting K to beyond 15 to 20 leads to minimal changes in the resulting inference.

A common technique in SMC is to resample the particles according to the importance weights $\{W_t^m\}_{m=1}^M$ when the effective sample size of the particles drops below a level. In sampling trees, however, the importance weights are affected by the choice of nodes to divide in multiple steps, and so the standard resampling scheme can be

too “short-sighted” and often results in sacrificing promising trees prematurely. To address this issue we follow the strategy proposed in Lu et al. (2013) by resampling the particles according to weights $a_t^m \propto (W_t^m)^\kappa$ for some $\kappa \in (0, 1]$, and compute the new importance weights proportional to W_t^m/a_t^m . We generally recommend using a moderate choice of κ such as 0.5, which we have found satisfactory in a variety of numerical experiments, and will be our default choice in all of our later examples.

3.2 Posterior computation given sampled tree structures

The second stage of our inference strategy is to compute the posterior distributions of the latent states $V(A)$ and the conditional probabilities $\theta_g(A)$ given each sampled tree. We shall focus on deriving a generic recipe for computing the posterior of the latent states given the tree that works for all models under consideration. Given both the tree and the latent states, the posterior of $\theta_g(A)$ boils down to the corresponding posterior of standard PT models on a dyadic tree, which is problem-specific as provided in the literature on each such model.

Now suppose the SMC algorithm has produced a collection of finite trees $\{T^m\}_{m=1}^M$ along with the importance weights $\{W^m\}_{m=1}^M$. Given each tree T^m , it is possible to analytically calculate the exact posterior of $\{V(A)\}_{A \in T^m}$ with recursive message passing (a form of dynamic programming), which we describe below.

For $A \in \mathcal{N}(T^m)$, let $\phi_A(i)$ be the marginal likelihood on A given that $V(A) = i$. (As before we suppress the dependency of the marginal likelihood on the data to simplify notation.) Specifically,

$$\phi_A(i) = \int q(\mathbf{x} | A) P(dq | V(A) = i), \text{ where } q(\mathbf{x} | A) = \prod_{g=1}^G \prod_{z \in x_g(A)} q_g(z | A). \quad (7)$$

In Eq. (7), taking the integration with respect to $P(dq | V(A) = i)$ is equivalent to integrating out $\theta_g(A)$ as well as the $\theta_g(A')$ and $V(A')$ terms for all descendants A' of A . Another useful quantity is the marginal likelihood on a node A given the state of

its parent node $V(A^p) = i$, which we denote as $\Phi_A(i)$ and is given by

$$\Phi_A(i) = \begin{cases} \prod_{x \in \mathbf{x}(A)} \mu(x | A) & \text{if } A \text{ is a leaf node,} \\ \sum_{i'=1}^I \xi_{i,i'}(A) \phi_A(i') & \text{if } A \text{ is a non-leaf node.} \end{cases} \quad (8)$$

Note that the $\Phi_A(i)$ and $\phi_A(i)$ terms are related by

$$\phi_A(i) = M_i(A | J(A)) \Phi_{A_l}(i) \Phi_{A_r}(i), \quad (9)$$

where M_i is the marginal likelihood defined in (4) given under the decision $J(A) = (D(A), L(A))$ to divide A into A_l and A_r . By iteratively computing Eqs. (8) and (9) in a bottom-up fashion (i.e., starting from the leaves all the way to the root), we can compute the pair $\{(\phi_A(i), \Phi_A(i)) : A, i\}$ for all nodes in the tree, and these pairs are the “messages” passed along the tree from leaf to root.

Given the values of $\{(\phi_A(i), \Phi_A(i)) : A, i\}$, we can now obtain the posterior Markov transition probability matrices of the latent states given the tree T , $\tilde{\xi}(A) = (\tilde{\xi}_{i,i'}(A))_{I \times I}$ where the (i, i') th element $\tilde{\xi}_{i,i'}(A) = P(V(\Omega) = i' | V(A^p) = i, \mathbf{x}, T)$ and the posterior marginal probabilities of the latent states given the tree T ,

$$\tilde{\gamma}(A) = (\tilde{\gamma}_i(A))_{i=1}^I = (P(V(\Omega) = i | \mathbf{x}, T))_{i=1}^I.$$

Specifically, by the Bayes’ theorem, $\tilde{\xi}(A) = D_1^{-1}(A) \xi(A) D_2(A)$, where $D_1(A)$ and $D_2(A)$ are diagonal matrices with $D_1(A)_{i,i} = \Phi_A(i)$ and $D_2(A)_{i,i} = \phi_A(i)$. After computing these transition matrices, we can compute $\tilde{\gamma}(A)$ (the feedback “message”) in the top-down manner (i.e., starting from the root and down to the leaves) as $\tilde{\gamma}(\Omega) = \tilde{\xi}_{1,\cdot}(\Omega)$ and $\tilde{\gamma}(A) = \tilde{\gamma}(A^p) \tilde{\xi}(A)$ for $A \neq \Omega$.

We note that $\Phi_\Omega(1)$ computed in the recursive algorithm is the overall marginal likelihood given the tree T , $P(\mathbf{x} | T)$, which can be used to find the *maximum a posteriori* (MAP) tree among the sampled trees, i.e., the sampled tree T^m that maximizes $P(T^m | \mathbf{x}) \propto P(T^m) P(\mathbf{x} | T^m)$. We can use this “representative” tree, along with the conditional posterior of the latent states given this tree, to visualize

and summarize the posterior inference in an interpretable way.

Now that we have completely described our inference algorithm, we provide two specific examples to demonstrate how one may use the output of the algorithm—namely the sampled trees along with the conditional posterior given the trees—to carry out inference. The first example is density estimation which involves learning the within-sample distributional features while the second is two-group comparison whose focus is on learning the between-sample structures. The inference strategies for these quintessential examples are generalizable to a variety of other tasks.

Example 1: Density estimation

We consider the problem of estimating an unknown density based on a set of i.i.d. observations from that density. This corresponds to $G = 1$ and so we can drop the subscript g to simplify the notation. We shall use the posterior mean density, also called the predictive density— $\mathbb{E}[q(\cdot) \mid \mathbf{x}]$ —as an estimate for the density $q = dQ/d\mu$. It can be computed by integrating out the random trees and the latent variables based on the SMC sample. Specifically, for a given tree T , we define for all $A \in T$ and $i \in \{1, 2, \dots, I\}$, the following quantity

$$e_A(i) := \mathbb{E}[Q(A)I_{\{V(A)=i\}} \mid \mathbf{x}].$$

All of these $e_A(i)$ terms can be computed together by a single top-down (i.e., root-to-leaf) recursion on the tree as given in the following proposition.

Proposition 3.2. *For the root node, $e_\Omega(i) = \tilde{\gamma}_{1,i}(\Omega)$. For a non-root node A , $e_A(i)$ can be computed recursively as*

$$e_A(i') = \sum_{i=1}^I \tilde{\xi}_{i,i'}(A) \mathbb{E}[\vartheta(A^p) \mid V(A^p) = i, T^m, \mathbf{x}] e_{A^p}(i), \text{ where}$$

$$\vartheta(A^p) = \begin{cases} \theta(A^p) & \text{if } A \text{ is the left child of } A^p, \\ 1 - \theta(A^p) & \text{if } A \text{ is the right child of } A^p. \end{cases}$$

Now with a recipe for obtaining all $e_i(A)$'s for a given tree, we can obtain the conditional predictive measure given the tree T ,

$$\mathbb{E}[Q(B) \mid \mathbf{x}, T] = \sum_{A \in \mathcal{L}(T)} \frac{\mu(B \cap A)}{\mu(A)} \sum_{i=1}^I e_A(i) \quad \text{for any Borel set } B.$$

Now, given an SMC sample of M trees and weights, the posterior predictive density can be computed by

$$\mathbb{E}[q(x) \mid \mathbf{x}] \approx \sum_{m=1}^M W^m \frac{\mathbb{E}[Q(B^m(x)) \mid \mathbf{x}, T^m]}{\mu(B^m(x))},$$

where $B^m(x) \in \mathcal{L}(T^m)$ the leaf node to which x belongs.

Example 2: Two-group comparison

If we are interested in comparing two groups of observations using generalizations to the PT models described in Section 2.3.1, we shall compute the posterior probability of the two hypotheses H_0 and H_1 . For example, when $V(A)$ is defined as in Eq (3), the posterior probability of the “global” null hypothesis $H_0 : Q_1 = Q_2$ is given by

$$\begin{aligned} P(H_0 \mid \mathbf{x}) &= \sum_{T \in \mathcal{T}} P(V(A) \neq 1 \text{ for all } A \in \mathcal{N}(T) \mid T, \mathbf{x}) P(T \mid \mathbf{x}) \\ &\approx \sum_{m=1}^M W^m P(V(A) \neq 1 \text{ for all } A \in \mathcal{N}(T^m) \mid T^m, \mathbf{x}), \end{aligned}$$

where the sum over \mathcal{T} in the first row is over all finite trees with maximum resolution K and the quantity $P(V(A) \neq 1 \text{ for all } A \in \mathcal{N}(T^m) \mid T^m, \mathbf{x})$ again is available analytically by message passing (details given in **Supplementary Materials B**).

In addition to testing the existence of any difference between groups, it is usually of interest to detect where and how the underlying distributions differ. To this end, we can compute the “posterior marginal alternative probability” (PMAP) on each node A , along any sampled tree T^m :

$$P(\theta_1(A) \neq \theta_2(A) \mid T^m, \mathbf{x}) = P(V(A) = 1 \mid T^m, \mathbf{x}) = \tilde{\gamma}_i(A).$$

Reporting the PMAPs along a representative tree such as the MAP among the sampled trees can be a particularly useful visualizing tool to help understand the nature of the underlying difference. One can also report on each A the estimated magnitude of the difference using a notion of “effect size” based on the log-odds ratio (Soriano and Ma, 2017), $\text{eff}(A) = \left| \log \left[\frac{\theta_1(A)}{1-\theta_1(A)} \right] - \log \left[\frac{\theta_2(A)}{1-\theta_2(A)} \right] \right|$. In particular, one can report the posterior expected effect size $\mathbb{E}[\text{eff}(A) \mid \mathbf{x}, T]$, which can be computed using a standard Monte Carlo (not MCMC) sample from the exact posterior given the representative tree. We will demonstrate this using a mass cytometry data set in Section 6

4 Theoretical Properties

Next we investigate the theoretical properties of the proposed model. Previous theoretical analysis on the PT had mostly focused on establishing the marginal posterior consistency and contraction of the random measures Q_g with respect to an unknown fixed truth (Walker and Hjort, 2001; Castillo, 2017b). We shall take a different perspective and instead investigate the asymptotic behavior of the marginal posterior of the partition tree T and the latent states as these are often critical quantities of practical importance in applications. We note that once given the tree and the latent states, our model reduces to standard PTs on a dyadic tree and thus the posterior consistency of the random measures Q_g ’s will follow from previous results once we establish the posterior consistency of the tree and the latent states. As the sample size increases, the two key theoretical questions of interest here are:

- (1) What tree structures does the marginal posterior of T concentrate around?
- (2) How does the posterior of the latent states given the tree behave?

These two questions have broad relevance in inference using PT models, and previously several authors have investigated the second question in the two-group comparison context for their variants of the PT model (Holmes et al., 2015; Soriano

and Ma, 2017). In addressing the second question more generally, we aim to provide results that encompass these previous analyses as special cases. According to our limited knowledge, we are not aware of previous studies on the first question.

We will address each of the two questions in turn. Throughout this section, we consider finite PTs with maximum depth of the trees set to some value K . We use \mathcal{T}^K to denote this collection of trees. Also, the asymptotic results are derived under the prior for $L(A)$ provided in Section 2.2 which can depend on the (finite) sample size. The case of an uniform priors on $L(A)$ independent of the sample size is included as a special case where the hyperparameter $\eta = 0$. Finally, we consider models that satisfy **Assumption 1** and **Assumption 2** described below. The models discussed in Section 2.3.1 all meet this requirement.

Assumption 1. *For each group $g \in \{1, \dots, G\}$, let n_g be the sample size and P_g the true probability measure from which the observations are generated. We assume that*

- (i) *There exists $\zeta_g \in (0, 1]$ such that $\zeta_g = \lim_{n \rightarrow \infty} \frac{n_g}{n}$ for $g \in \{1, \dots, G\}$, where $n = n_1 + \dots + n_G$ is the total number of observations across all groups.*
- (ii) *The sampling distribution P_g satisfies $P_g \ll \mu$, and the density $p_g = dP_g/d\mu$ is positive almost everywhere.*

Additionally, given the tree T and the latent states, the parameters $\{\theta_g(A)\}_{g=1}^G$ are given one of the following priors (the model can adopt a mix of these priors for different combinations of A and $V(A)$ values):

Prior A : $\theta_g(A)$ *independently follow a beta prior.*

Prior B : $\theta_1(A) = \dots = \theta_G(A)$ *and follow a beta prior.*

Prior C : $\theta_1(A) = \dots = \theta_G(A) \equiv c(A)$, *some constant in $(0, 1)$.*

Establishing the theoretical properties also requires a condition on the latent states. In particular, under some states, the support of the prior of the parameters $\{\theta_g(A)\}_{g=1}^G$ needs to include the true conditional probabilities. To describe this

requirement, given a tree $T \in \mathcal{T}^K$, let $S_i(A | T)$ be the support of the prior on $(\theta_1(A), \dots, \theta_G(A))$ under the state $V(A) = i$. Then, let $\tau(A | T)$ denote the collection of “feasible states” on A . (A state is “feasible” if the true conditional probabilities are in the support of the corresponding prior given the state.) That is,

$$\tau(A | T) := \{i \in \{1, \dots, I\} : (P_1(A_l | A), \dots, P_G(A_l | A)) \in S_i(A | T)\}.$$

The next assumption states that the prior for the latent states must give positive probability for all the latent states to all simultaneously be feasible.

Assumption 2. *For every $T \in \mathcal{T}^K$, $P(V(A) \in \tau(A | T) \text{ for all } A \in T) > 0$.*

With these assumptions, we next derive asymptotic properties for the marginal posteriors for the tree and the state variables. In the following, we use the notation \mathbf{x}_n instead of \mathbf{x} for the data to indicate the total sample size.

In order to describe the posterior convergence of the partition trees, we introduce a notion for “tree-based approximation for probability measures”. Let T be a finite tree and H a probability measure. Then the “tree-based approximation of H under T ”, denoted by $H|_T$, is defined as $H|_T(B) = \sum_{A \in \mathcal{L}(T)} H(A) \frac{\mu(B \cap A)}{\mu(A)}$, for any $B \in \mathcal{B}(\Omega)$. The following theorem then characterizes the trees the posterior concentrates on as the sample size grows.

Theorem 4.1. *Let \mathcal{T}_M^K be the collection of trees under which the tree-based approximation of the measures P_g minimizes the Kullback-Leibler divergence from the P_g ’s plus a penalty term on unbalanced splits. That is,*

$$\mathcal{T}_M^K = \arg \min_{T \in \mathcal{T}^K} \sum_{g=1}^G \zeta_g \{ \text{KL}(P_g || P_g|_T) + \eta B_g(T) \}, \quad (10)$$

where

$$B_g(T) = \sum_{A \in \mathcal{N}(T)} P_g(A) f \left(\left| \frac{\mu(A_l)}{\mu(A)} - 0.5 \right| \right).$$

Then the marginal posterior of T concentrates on \mathcal{T}_M^K . That is, as $n \rightarrow \infty$,

$$P(T \in \mathcal{T}_M^K | \mathbf{x}_n) \xrightarrow{p} 1.$$

For the state variables, it is desirable that their posterior distribution concentrates on a collection of feasible states. Moreover, when multiple configurations of the states are feasible, it is desirable that the posterior concentrates around such configurations that provide the most parsimonious representation of the true distributions. For example, if the true conditional distribution on a node is uniform, a model that introduces a possible non-uniform structure on this node is feasible but redundant. White and Ghosal (2011) and Li and Ghosal (2014) showed that, in quite general settings of multi-resolution inference, the posterior probability of such redundant models tends to concentrate its mass on 0. By adapting their techniques, we show that the same property holds in the case of our model.

To formally describe the results, we need to define the complexity of the model specified by the latent states. Given the state $V(A) = i$, the complexity of the $\{\theta_g(A)\}_{g=1}^G$, in other words, the number of free parameters of the prior distribution under the i th state is denoted by $C_i(A)$. For example, for two-group comparison,

$$C_i(A) = \begin{cases} 1 & \text{if } \theta_1(A) = \theta_2(A) \\ 2 & \text{if } \theta_1(A) \neq \theta_2(A). \end{cases}$$

Next we introduce the complexity of a combination of states on the tree T . Given a tree T , let \mathbf{V} denote a combination of the state variables $\{V(A)\}_{A \in \mathcal{N}(T)}$ and let $\mathbf{v} = \{\mathbf{v}(A)\}_{A \in \mathcal{N}(T)}$ ($\mathbf{v}(A) \in \{1, \dots, I\}$) be one of the possible realizations of \mathbf{V} . Then we define the model complexity under \mathbf{v} as follows:

$$C(\mathbf{v}) = \sum_{A \in \mathcal{N}(T)} C_{\mathbf{v}(A)}(A). \quad (11)$$

The next theorem shows that the posterior distribution of the states given the tree will concentrate on those that are feasible and most parsimonious.

Theorem 4.2. *For $T \in \mathcal{T}^K$, let $\mathcal{V}_T = \{\mathbf{v} : \mathbf{v}(A) \in \tau(A \mid T) \text{ for all } A \in \mathcal{N}(T)\}$. Then $P(\{\mathbf{V} \in \mathcal{V}_T\} \cap \{C(\mathbf{V}) = \min_{\mathbf{v} \in \mathcal{V}_T} C(\mathbf{v})\} \mid T, \mathbf{x}_n) \xrightarrow{P} 1$.*

Remark: Consistency results for several existing models are special cases of this

theorem. For example, we derive the consistency of the MRS model for two-group comparison as a corollary in **Supplementary Materials C**.

5 Experiments

In this section, we carry out simulation studies to examine the performance of our model and inference algorithm. In particular, we are interested in (i) understanding how the model with the flexible tree prior compares to those with a “divide-in-the-middle” restriction, and (ii) verifying the linear scalability of our inference algorithm with respect to increasing dimensionality. We again consider the two quintessential examples—(i) density estimation and (ii) the two-group comparison—for inferring within-group and between-group structures respectively.

We shall consider both low-dimensional settings where the underlying structure is easy to interpret and software for existing PT models are available, and high-dimensional settings for which existing implementation of PT models is not applicable and we use our SMC algorithm to carry out inference for both our model and the earlier models with fixed partitioning points (which are special cases of our model). Throughout the experiments, the parameters N_L and M are fixed to 32 and 1000 respectively. We note that larger N_L values can also be adopted at a linear computational cost but did not lead to noticeable change in performance in our examples. Details such as the settings of hyper-parameters and simulated data sets are provided in **Supplementary Materials F** unless explicitly described in this section.

5.1 Density estimation

We first consider 2D examples to observe what kind of tree structures are obtained under the flexible model and how prior specification in Eq. (2) influences the performance. After that, we move to higher dimensional cases to examine the scalability of our new SMC method and the effect of incorporating the flexible partition. For this

task we compare our model with the APT model (Ma, 2017) which also incorporates a prior on the dimension to divide but restricts partitioning at middle points. Its posterior computation is implemented by the `apt` function in the R package `PTT`.

5.1.1 Two-dimensional cases

Simulated data are generated from the three scenarios with the densities visualized in the first row of Figure 3. (Details on the simulation settings are provided **Supplementary Materials F.1.2.**) Also presented in Figure 3 are examples of the posterior mean densities $\mathbb{E}[q \mid \mathbf{x}]$ as well as the partition blocks under the MAP tree. Note that the posterior mean is computed by integrating out the unknown tree, and the MAP tree is presented to visualize key distributional features. The results for the first scenario confirms that our more flexible model is much more effective in capturing the discontinuous boundaries of the true density. For the second scenario, our model tends to draw the boundaries that surround the true clusters. In the trees given under the different values of η , however, we can see that fewer nodes were divided inside the clusters when $\eta = 0.01$. In contrast, when $\eta = 0.1$, the representative tree draws outlines of the clusters and divides regions inside of the clusters at the same time. A similar phenomenon is observed in the third scenario—under our model with flexible partitioning points, partition lines are formed around the region with high density, when $\eta = 0.1$ for the boundaries were also drawn within the high probability region. The quantitative comparison based on the KL divergence is provided in **Supplementary Materials G**, which is consistent with the explanation above.

5.1.2 Higher-dimensional cases

We generate d -dimensional i.i.d. observations from a density with independent pairs of margins, i.e., $f(x_1, x_2, \dots, x_d) = \prod_{j=1}^{d/2} f_j(x_{2j-1}, x_{2j})$ where

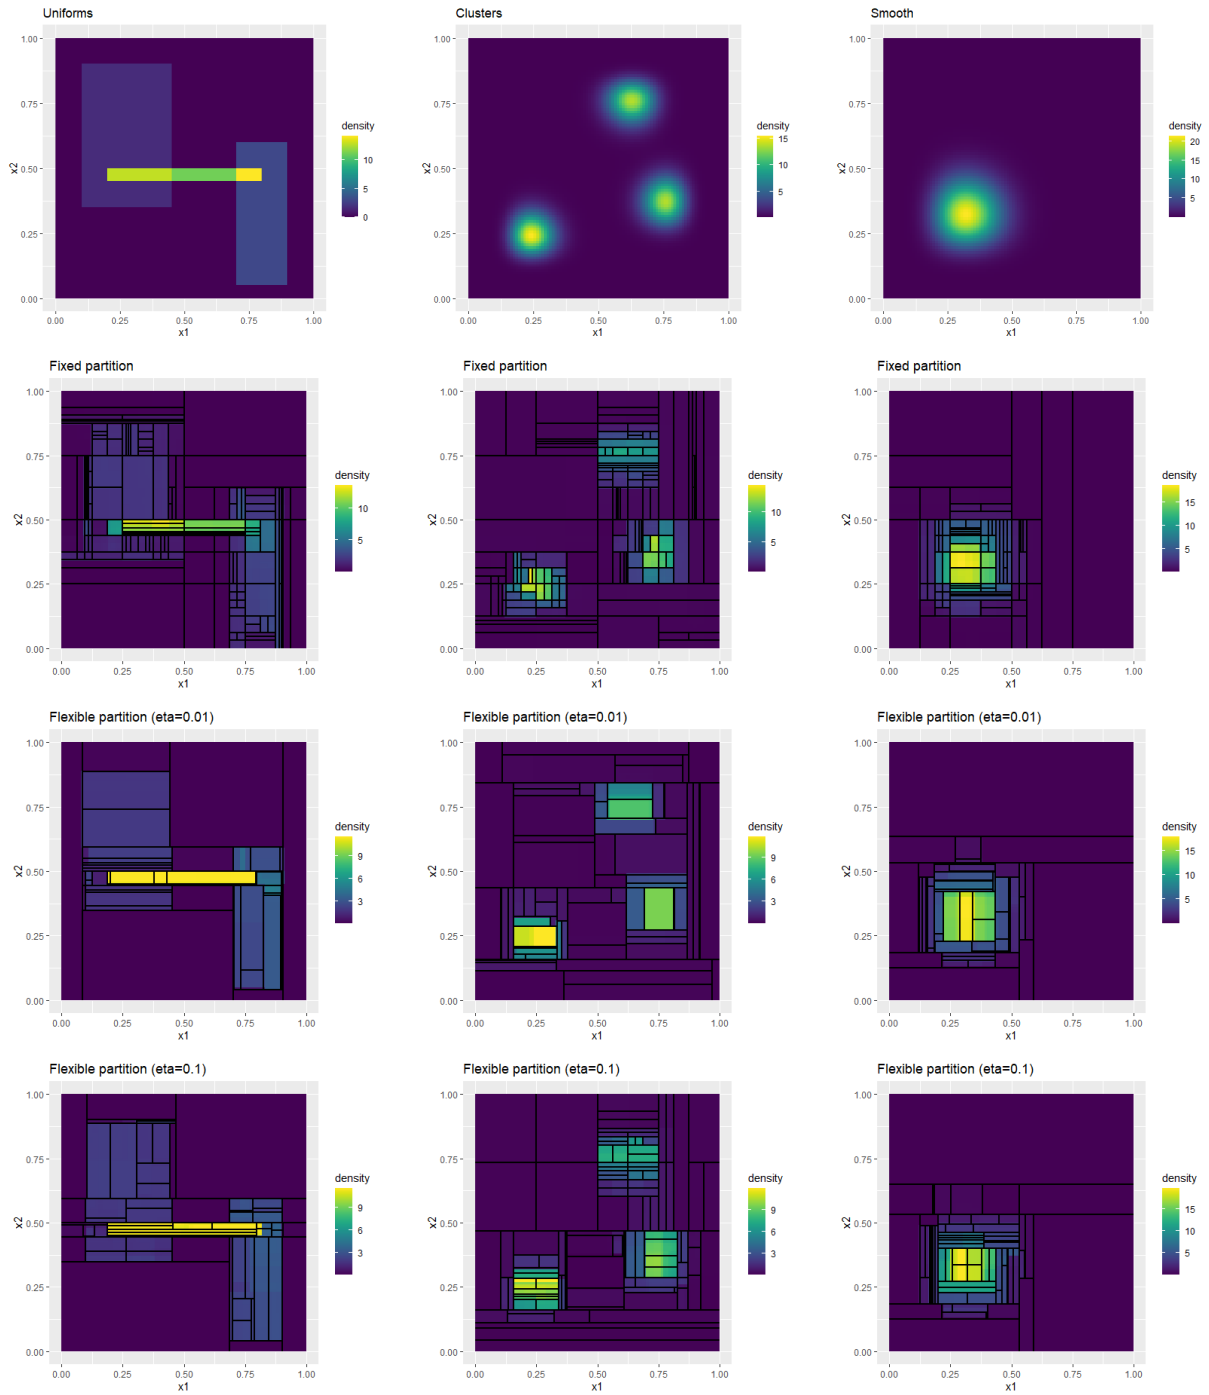


Figure 3: The posterior means of the densities and the representative trees obtained under $n = 1000$. Each column corresponds to a simulation scenario. The first row shows that true densities, the second row corresponds to the APT model (with fixed partition), and the third and fourth rows correspond to our model with flexible partitioning with parameters $\eta = 0.01$ and 0.1 respectively.

$$f_j(x_{2j-1}, x_{2j}) = p_j \text{Beta}(x_{2j-1} \mid 0.25, 1) \times \text{Beta}(x_{2j} \mid 0.25, 1) \\ + (1 - p_j) \text{Beta}(x_{2j-1} \mid 50/j, 50/j) \times \text{Beta}(x_{2j} \mid 50/j, 50/j),$$

with $p_j = 0.25 + 0.7/j$. We consider two different situations: (i) the dimension $d = 6$, and the sample size n changes from 5,000 to 50,000; and (ii) the sample size $n = 10,000$ and the dimensionality changes from 10 to 100. For our method, the maximum depth K is set to 15. For the algorithm of the original APT (implemented by the `apt` function in the `PTT` package), in the first case with $d = 6$, the maximum resolution is fixed to 9 because setting higher values leads to insufficient memory. Also, because the `apt` function does not scale if the dimension is beyond $d \approx 10$, in the second case with large d , we used the proposed SMC algorithm to carry out inference for the original APT model, which corresponds to setting $N_L = 2$, namely, fixing all the partition points to the middle. In addition to the original APT, we show the comparison with the classical PT method in which each node is split into 2^d and the maximum depth is 15, the Dirichlet process Gaussian mixture model (Escobar and West, 1995; Müller et al., 1996) implemented by the `PYdensity` function in the R package `BNPmix` (Corradin et al., 2021), and for (i) the Gaussian kernel density estimation using the `kde` function in the R package `ks` (Duong et al., 2007).




Figure 4: The wall time under five different data sets. The flexible model with $\eta = 0.01$ is used. The black dashed lines indicate the average times.

Figure 4 presents the computational time for five different data sets. To obtain the

result, we used a single-core environment using Intel Xeon Gold 6154 (3.00 GHz) CPU. The computational time is linear in both the sample size and the dimensionality.

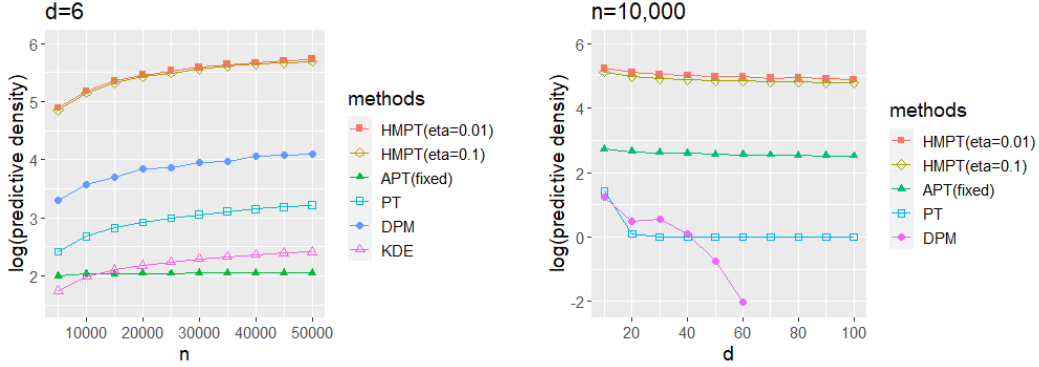


Figure 5: Predictive performance of 6 methods. Each point corresponds to the average of the predictive score in Eq. (12) based on 50 data sets. Each interval is formed by adding and subtracting the standard deviation. In the right plot, the predictive scores of the DPM model for the over 60 dimensional cases are below the displayed range.

Because in the high-dimensional settings we cannot obtain the KL divergence between the estimated density and the true density through numerical approximation, we compare the models based on predictive scores, which is in essence an empirical estimate of the KL divergence. Specifically, for each training set \mathbf{x} , we can generate a new test set denoted by $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ from the same true model and compute

$$\frac{1}{n} \sum_{i=1}^n \log \hat{p}(x_i^* | \mathbf{x}), \quad (12)$$

where \hat{p} is the posterior mean of the density $\mathbb{E}[q | \mathbf{x}]$. We repeat this computation for 50 test/training set pairs and take the average. The results, given in Figure 5, show that our model substantially outperforms the competitors by this criteria both when $d = 6$ with varying sample size and when n is fixed with varying dimensionality. We also investigate the performance under small sample sizes, and the results are similar. (See **Supplementary Materials G** for details.)

5.2 Two-group comparison

Next we consider the two-group comparison problem, evaluate the performance of our model, and compare it to the original MRS with the “divide-in-the-middle” restriction. We use three scenarios (“Local location shift”, “Local dispersion difference”, and “Correlation”) to generate 50-dimensional data sets. (Details of the scenarios are provided in **Supplementary Materials F.2.2**.) The first two scenarios involve two-group difference that lies in only parts of the sample space, or “local” differences, which will help demonstrate the usefulness of inferring the partition tree in identifying the nature of the differences. The sample size is $n_1 = n_2 = 2,000$ in all scenarios.

The original algorithm for inference under the MRS model by message passing, which is implemented by the `mrs` function in the R package `MRS`, is not scalable beyond about 10 dimensions even with fixed partition locations. Hence we compute the posterior for both our model and the original MRS in all scenarios with our SMC and message passing hybrid algorithm. For the two models, the maximum resolution K is fixed to 15. We compare the performance using receiver operating characteristic (ROC) curves computed based on 200 simulated data sets under each scenario.

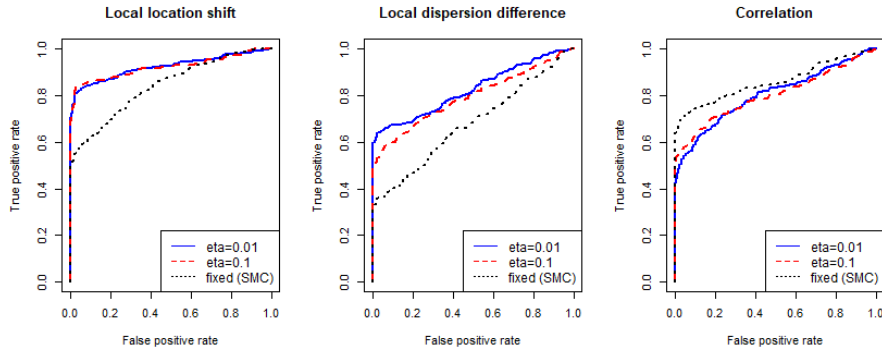


Figure 6: The ROC curves for the 50-dimensional examples.

Figure 6 presents the ROC curves. For the location shift and dispersion differences, the model with flexible partitioning results in substantially higher sensitivity. For the




Figure 7: The node with the highest PMAP $P(V(A) = 1 \mid \mathbf{x})$ under the three scenarios for the 50-dimensional example, estimated by the MRS with flexible partitioning and $\eta = 0.1$. The solid lines mark the boundaries of the nodes and the partition line that divides them into the two children nodes. The red triangle points and the blue circle points are the observations of the two groups included in the node. Gray points indicate the observations outside the node.

correlation scenario, the model with fixed partitioning locations performed slightly better. This is not surprising since in this scenario the difference exists smoothly over entire ranges of the dimensions without natural “optimal” division points, and so the performance gap is the cost for searching over more possible partition locations, none of which improves the model fit than the middle point. It is worth noting again that while the model with fixed partitioning performs well here, it is only with our new computational algorithm that it can be fit to data of such dimensionality.

To demonstrate the posterior model can help understand the nature of the differences, we present under each scenario the node with the highest PMAP, or $P(V(A) = 1 \mid \mathbf{x}) = P(\theta_1(A) \neq \theta_2(A) \mid \mathbf{x})$, in Figure 7. In the location shift and dispersion difference scenarios the boundaries are away from the middle point to characterize the difference, which partly explains the sensitivity gain in adopting the flexible tree prior.

6 Application to a mass cytometry data set

Finally, we apply our model for two-group comparison to a mass cytometry data set collected by Kleinsteuber et al. (2016). The data set records 19 different mea-

measurements including physical measurements and biomarkers on single cells in blood samples from a group of HIV patients as well as in reference samples from healthy donors. For demonstration, we compare the sample from an individual patient sample (Patient #1) and to that from a healthy donor to identify differences in immune cell profiles from these samples. The sample sizes are 29,226 for the health donor and 228,498 for the patient, with each observation corresponding to a cell. We set $\eta = 0.1$ and the maximum depth K to 25.

Given the large sample sizes, the posterior probability for the global alternative $P(Q_1 \neq Q_2 \mid \mathbf{x})$ is almost 1 and so is of less interest. Our focus is instead on identifying the cell subsets on which the samples differ and on quantifying such differences. To this end, we identify a representative tree and report the “effect size” (i.e., the posterior expected log-odds ratio between the two samples) on each node in a representative tree—the MAP among the sampled trees.

The estimated $\text{eff}(A)$ ’s on the MAP tree up to level 9 is visualized in Figure 8, and the full tree is provided in **Supplementary Materials G**. We note that the nodes on which there is significant evidence for two-group differences, as well as those with large estimated effect sizes tend to be nested or clustered in subbranches of the tree, which is consistent with our intuition that there is spatial correlation in the two-group differences, and justifies the hidden Markov structure embedded in the MRS model.

Figure 14 in **Supplementary Materials G** presents the 20 nodes with the largest values of estimated $\text{eff}(A)$. In this figure, many of the nodes are in very deep levels of the tree. We adopted a spike-and-slab with higher spike probability in very deep tree levels to further speed up the computation (details given in **Supplementary Materials A**) and that explains why many of the very deep, small nodes plotted have partition lines in the middle under the MAP tree.




Figure 8: The MAP tree for the mass cytometry data set visualized up to the 9th level. The size and the color indicate the estimated $\text{eff}(A)$, and the number above a node indicates dimension in which it is split. Only the nodes with the sample size larger than 50 are drawn.

7 Concluding Remarks

We have proposed a general framework for the PT model that incorporates a flexible prior on the partition tree and can accommodate latent state variables with Markov dependency along the partition tree. We have proposed a sampling algorithm that combines SMC and recursive message passing that can scale up to moderately high-dimensional (~ 100 -dim) problems. Our numerical experiments confirm that our sampling algorithm scales linearly in the sample size and the flexible partitioning tree prior can result in substantial gain in performance in some settings. Though we have mainly used two inference tasks—namely density estimation and two-group comparison—to demonstrate our model and algorithm, our approach can be readily applied to other PT models with a hidden Markov structure.

One notable limitation of our model—and in fact all CART-like models—is that we only consider trees in which the node boundaries are all parallel to the axes. This could lead to inefficiency in inference. For example, when there is a strong correlation between several variables, drawing boundaries slanted according to the correlation structure would be more effective in characterizing the underlying distribution. Such

trees will need to be represented by more than just the $D(A)$ and $L(A)$ used in our model, and how to efficiently compute the posterior distribution is an open problem.

Finally, our proposed algorithm is currently designed to be run in a single computer environment, so though the computational cost is linear to the sample size n , dealing with huge data with huge n (e.g., $> 10^{12}$) is not yet realistic. It is of future interest to develop parallelized versions of the algorithm for distributed systems, which could explore either the parallel structure over nodes or parallel SMC algorithms.

Software

An R package for our method is available at <https://github.com/MaStatLab/SMCMP>.

Acknowledgment

LM's research is partly supported by NSF grants DMS-2013930 and DMS-1749789.

NA is partly supported by a fellowship from the Nakajima Foundation.

References

- Berger, J. O. and Guglielmi, A. (2001). Bayesian and conditional frequentist testing of a parametric model versus nonparametric alternatives. *Journal of the American Statistical Association*, 96(453):174–184.
- Castillo, I. (2017a). Pólya tree posterior distributions on densities. In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, volume 53, pages 2074–2102. Institut Henri Poincaré.
- Castillo, I. (2017b). Pólya tree posterior distributions on densities. In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, volume 53, pages 2074–2102. Institut Henri Poincaré.

- Castillo, I. and Randrianarisoa, T. (2021). Optional Pólya trees: posterior rates and uncertainty quantification. *arXiv preprint arXiv:2110.05265*.
- Chen, Y. and Hanson, T. E. (2014). Bayesian nonparametric k-sample tests for censored and uncensored data. *Computational Statistics & Data Analysis*, 71:335–346.
- Chipman, H. A., George, E. I., and McCulloch, R. E. (1998). Bayesian cart model search. *Journal of the American Statistical Association*, 93(443):935–948.
- Christensen, J. and Ma, L. (2020). A Bayesian hierarchical model for related densities by using Pólya trees. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1):127–153.
- Corradin, R., Canale, A., and Nipoti, B. (2021). Bnpmix: an r package for bayesian nonparametric modelling via pitman-yor mixtures. *Journal of Statistical Software*.
- Crouse, M. S. and Baraniuk, R. G. (1997). Contextual hidden markov models for wavelet-domain signal processing. In *Conference Record of the Thirty-First Asilomar Conference on Signals, Systems and Computers*, volume 1, pages 95–100.
- Duong, T. et al. (2007). ks: Kernel density estimation and kernel discriminant analysis for multivariate data in r. *Journal of Statistical Software*, 21(7):1–16.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*, 90(430):577–588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of statistics*, pages 209–230.
- Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *Annals of statistics*, 2(4):615–629.
- Filippi, S., Holmes, C. C., et al. (2017). A Bayesian nonparametric approach to testing for dependence between random variables. *Bayesian Analysis*, 12(4):919–938.
- Freedman, D. A. (1963). On the asymptotic behavior of bayes’ estimates in the discrete case. *Annals of Mathematical Statistics*, pages 1386–1403.

- Hanson, T. and Johnson, W. O. (2002). Modeling regression error with a mixture of Polya trees. *Journal of the American Statistical Association*, 97(460):1020–1033.
- Hanson, T. E. (2006). Inference for mixtures of finite Polya tree models. *Journal of the American Statistical Association*, 101(476):1548–1565.
- Holmes, C. C., Caron, F., Griffin, J. E., Stephens, D. A., et al. (2015). Two-sample Bayesian nonparametric hypothesis testing. *Bayesian Analysis*, 10(2):297–320.
- Jara, A. and Hanson, T. E. (2011). A class of mixtures of dependent tail-free processes. *Biometrika*, 98(3):553–566.
- Kleinsteuber, K., Corleis, B., Rashidi, N., Nchinda, N., Lisanti, A., Cho, J. L., Medoff, B. D., Kwon, D., and Walker, B. D. (2016). Standardization and quality control for high-dimensional mass cytometry studies of human samples. *Cytometry Part A*, 89(10):903–913.
- Lakshminarayanan, B., Roy, D., and Teh, Y. W. (2013). Top-down particle filtering for Bayesian decision trees. In *International Conference on Machine Learning*, pages 280–288.
- Lavine, M. (1992). Some aspects of Polya tree distributions for statistical modelling. *Annals of statistics*, 20(3):1222–1235.
- Li, M. and Ghosal, S. (2014). Bayesian multiscale smoothing of gaussian noised images. *Bayesian analysis*, 9(3):733–758.
- Lu, L., Jiang, H., and Wong, W. H. (2013). Multivariate density estimation by Bayesian sequential partitioning. *Journal of the American Statistical Association*, 108(504):1402–1410.
- Ma, L. (2017). Adaptive shrinkage in Pólya tree type models. *Bayesian Analysis*, 12(3):779–805.
- Ma, L. and Soriano, J. (2018). Analysis of distributional variation through graphical multi-scale beta-binomial models. *Journal of Computational and Graphical Statistics*, 27(3):529–541.

- Ma, L. and Wong, W. H. (2011). Coupling optional Pólya trees and the two sample problem. *Journal of the American Statistical Association*, 106(496):1553–1565.
- Muliere, P. and Walker, S. (1997). A Bayesian non-parametric approach to survival analysis using Polya trees. *Scandinavian Journal of Statistics*, 24(3):331–340.
- Müller, P., Erkanli, A., and West, M. (1996). Bayesian curve fitting using multivariate normal mixtures. *Biometrika*, 83(1):67–79.
- Neath, A. A. (2003). Polya tree distributions for statistical modeling of censored data. *Advances in Decision Sciences*, 7(3):175–186.
- Nieto-Barajas, L. E. and Müller, P. (2012). Rubbery Polya tree. *Scandinavian Journal of Statistics*, 39(1):166–184.
- Paddock, S. M. (2002). Bayesian nonparametric multiple imputation of partially observed data with ignorable nonresponse. *Biometrika*, 89(3):529–538.
- Paddock, S. M., Ruggeri, F., Lavine, M., and West, M. (2003). Randomized Polya tree models for nonparametric Bayesian inference. *Statistica Sinica*, pages 443–460.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of statistics*, 6(2):461–464.
- Soriano, J. and Ma, L. (2017). Probabilistic multi-resolution scanning for two-sample differences. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(2):547–572.
- Walker, S. and Hjort, N. L. (2001). On Bayesian consistency. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(4):811–821.
- Walker, S. G., Damien, P., Laud, P. W., and Smith, A. F. (1999). Bayesian non-parametric inference for random distributions and related functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):485–527.
- Walker, S. G. and Mallick, B. K. (1997). Hierarchical generalized linear models and frailty models with Bayesian nonparametric mixing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(4):845–860.
- White, J. T. and Ghosal, S. (2011). Bayesian smoothing of photon-limited images

with applications in astronomy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):579–599.

Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Annals of Mathematical Statistics*, 9(1):60–62.

Wong, W. H. and Ma, L. (2010). Optional Pólya tree and Bayesian inference. *Annals of Statistics*, 38(3):1433–1459.

A Spike-and-slab type prior for $L(A)$

A.1 Introducing an auxiliary variable

The location variable $L(A)$ follows a spike-and-slab type prior which is expressed with an auxiliary variable $R(A)$ as

$$L(A) \mid R(A) \sim \mathbf{1}_{\{R(A)=1\}}\delta_{1/2}(\cdot) + \mathbf{1}_{\{R(A)=0\}} \sum_{l=1, l/N_L \neq 1/2}^{N_L-1} \hat{\beta}_l(A)\delta_{l/N_L}(\cdot),$$

where $\mathbf{1}$ is the indicator function and the sum of the parameters $\beta_l(A)$ is 1. Under this prior, $L(A)$ follows the prior degenerated at $1/2$ if $R(A) = 1$ and otherwise follows the distribution on grid points other than the middle point. $R(A)$ follows an asymmetric hidden Markov process

$$P(R(A) = 1 \mid R(A^p) = 0) = r(A),$$

$$P(R(A) = 1 \mid R(A^p) = 1) = 1,$$

where $r(A) \in [0, 1]$. $R(A) = 1$ is the absorbing state, so once A is divided at the middle point, $L(A') = 1/2$ for every A 's descendant node A' . In the estimation we especially set the parameters as follows:

$$r(A) = \beta_{N_L/2}, \quad \hat{\beta}_l = \frac{\beta_l(A)}{\sum_{l=1, l/N_L \neq 1/2}^{N_L-1} \beta_l(A)},$$

where $\{\beta_l(A)\}_{l=1, \dots, N_L}$ is given in (2). Under this setting the prior of $L(A)$ satisfies

$$P(L(A) = l \mid R(A^p) = 0) = \beta_l(A), \quad l = 1, \dots, N_L - 1.$$

Hence, $L(A)$ follows the same prior as defined in (1) unless A 's parent node is divided at the middle point, so the spike-and-slab prior can be seen as a natural extension.

A.2 SMC algorithm

In the SMC algorithm, we sample values of $R(A)$ in addition to $D(A)$ and $L(A)$. If $R(A) = 1$, which is equivalent to $L(A) = 1/2$, is sampled, we conclude there is no interesting structure on the node A so fix L to $1/2$ for all the subsequent nodes. Hence, we need to generalize the SMC algorithm discussed in Section 3.1 to sample from the joint posterior distribution of the finite trees and the auxiliary variables R .

To describe this joint posterior, let T_t denote the finite tree structure, which is determined by the sequence of decisions $J_{1:t}$ dividing the nodes $A_{1:t}$, and let $R_{1:t}$ be a sequence of the re-fixing variables for $A_{1:t}$. Then the target distribution we want to sample from in the SMC is defined as

$$\pi_t(T_t, R_{1:t}) = P(T_t, R_{1:t} \mid \mathbf{x}) \propto P(T_t, R_{1:t})P(\mathbf{x} \mid T_t, R_{1:t}).$$

The prior $P(T_t, R_{1:t}) = P(J_{1:t}, R_{1:t})$ have a Markov chain structure on the tree, and its transition probability is decomposed as

$$P(J_t, R_t \mid J_t^p, R_t^p) = P(R_t \mid R_t^p)P(J_t).$$

where $J_t^p = J(A_t^p)$ and $R_t^p = R(A_t^p)$ (A_t^p is the parent node of A_t). On the other hand, because $R_{1:t}$ are conditionally independent of the observations given T_t , the likelihood only depends on T_t as follows:

$$P(\mathbf{x} \mid T_t, R_{1:t}) = P(\mathbf{x} \mid T_t).$$

The likelihood has the same form as in the original case without the auxiliary variables R 's. Thus we can obtain the following proposition as a generalization of Proposition 3.1.

Proposition A.1. *Let $h(J_t | A_t)$ be a function of J_t defined as*

$$h(J_t | A_t) = \sum_{i=1}^I \left\{ \sum_{j=1}^I \varphi_j(A_t^p) \xi_{j,i}(A_t) \right\} M_i(A_t | J_t) \frac{\mu(\mathbf{x}(A_{t,l}) | A_{t,l}) \mu(\mathbf{x}(A_{t,r}) | A_{t,r})}{\mu(\mathbf{x}(A_t) | A_t)}.$$

Then the target distribution $\pi_t(T_t, R_{1:t})$ is expressed with $\pi_t(T_{t-1}, R_{1:t-1})$ as

$$\pi_t(T_t, R_t) = C \pi_t(T_{t-1}, R_{1:t}) \pi_t(R_t | T_{t-1}, R_t^p) \pi_t(J_t | T_{t-1}, R_t) w_t(T_{t-1}, R_t^p),$$

where C is a constant and

$$\begin{aligned} \pi_t(R_t | T_{t-1}, R_t^p) &= \frac{P(R_t | R_t^p) \sum_{j_t} P(j_t | R_t) h(j_t | A_t)}{\sum_{i=0,1} P(R_t = i | R_t^p) \left\{ \sum_{j_t} P(j_t | R_t = i) h(j_t | A_t) \right\}}, \\ \pi_t(J_t | T_{t-1}, R_t) &= \frac{P(J_t | R_t) h(J_t | A_t)}{\sum_{j_t} P(j_t | R_t) h(j_t | A_t)}, \\ w_t(T_{t-1}, R_t^p) &= \sum_{i=0,1} P(R_t = i | R_t^p) \left\{ \sum_{j_t} P(j_t | R_t = i) h(j_t | A_t) \right\}. \end{aligned}$$

The summation with j_t is taken over all possible decisions.

Its proof is essentially the same as Proposition 3.1 and Corollary 3.1 so it is omitted in this material. The conditional posteriors $\pi_t(R_t | R_t^p, T_{t-1})$ and $\pi_t(J_t | R_t, T_{t-1})$ are analytically obtained as follows. First, if $R_t^p = 0$, $\pi_t(R_t | R_t^p, T_{t-1})$ is *Bernoulli*($\tilde{r}(A_t)$),

where

$$\begin{aligned} \tilde{r}(A_t) &= r(A_t) \sum_{j=1}^J \lambda_j(A_t) h((j, 1/2) | A_t) \\ &\times \left[r(A_t) \sum_{j=1}^J \lambda_j(A_t) h((j, 1/2) | A_t) + (1 - r(A_t)) \sum_{j=1}^J \sum_{l=1}^{N_L-1} \lambda_j(A_t) \beta_l(A_t) h((j, l/N_L) | A_t) \right]^{-1}, \end{aligned}$$

If $R_t^p = 1$, then R_t is fixed to 1. Second, if $R_t = 0$, the posterior of D_t and L_t is the same distribution given in Section 3.1. On the other hand, if $R_t = 1$, L_t is fixed to 1/2, and the conditional posterior $\pi_t(J_t | L_t, R_t, T_{t-1})$ is $\text{Mult}(\hat{\lambda}_1(A_t), \dots, \hat{\lambda}_d(A_t))$, where

$$\hat{\lambda}_j(A_t) \propto \lambda_j(A_t) h(j, 1/2 | A_t).$$

After sampling (R_t, J_t) , the incremental weight $w_t(T_{t-1}, R_t^p)$ ($R_t^p = 0, 1$) is computed as

$$\begin{aligned} w_t(T_{t-1}, 0) &= r(A_t) \sum_{j=1}^J \lambda_j(A_t) h((j, 1/2) | A_t) + (1 - r(A_t)) \sum_{j=1}^J \sum_{l=1}^{N_L-1} \lambda_j(A_t) \beta_l(A_t) h((j, l/N_L) | A_t), \\ w_t(T_{t-1}, 1) &= \sum_{j=1}^J \lambda_j(A_t) h((j, 1/2) | A_t), \end{aligned}$$

with which we update the importance weight W_t as $W_t \propto W_{t-1} w_t(T_{t-1}, R_t^p)$.

The procedure to update the particle system $\{T_{t-1}^m, W_{t-1}^m\}_{m=1}^M$ to obtain $\{T_t^m, W_t^m\}_{m=1}^M$ is described in the following algorithm. The operations involving the index m is repeated for $m = 1, \dots, M$.

1. Choosing the current node

From T_{t-1}^m , choose the oldest node from the leaf nodes, which is denoted by A_t .

2. Obtaining the information of the parent node

Pick up A_t 's parent node, which is denoted by A_t^p , and load the values of $\varphi_i(A_t^p)$ for $i = 1, \dots, I$ and $R_t^{m,p} = R(A_t^p)$.

3. Computing the necessary quantities

If $R_t^{m,p} = 0$, compute $M_i(A \mid j, l/N_L)$ ($i = 1, \dots, I$) and $h(j, l/N_L \mid A_t)$ for $j = 1, \dots, d$ and $l = 1, \dots, N_L - 1$.

If $R_t^{m,p} = 1$, compute $M_i(A \mid j, 1/2)$ ($i = 1, \dots, I$) and $h(j, 1/2 \mid A_t)$ for $j = 1, \dots, d$.

4. Deciding whether to fix the partition or not

If $R_t^{m,p} = 0$, compute the parameter $\tilde{r}(A_t)$ and draw $R_t^m \sim \text{Bernoulli}(\tilde{r}(A_t))$.

If $R_t^{m,p} = 1$, set R_t^m to be 1.

5. Dividing the current node

Sample $J_t^m = (D_t^m, L_t^m)$ as follows:

- If $R_t^m = 0$, compute the parameters $\tilde{\lambda}_j(A_t)$ for $j = 1, \dots, d$ and sample

$$D_t^m \sim \text{Mult}(\tilde{\lambda}_1(A_t), \dots, \tilde{\lambda}_d(A_t)).$$

Given D_t^m , compute the parameters $\tilde{\beta}_l(A_t)$ for $l = 1, \dots, N_L - 1$ and sample

$$L_t^m \sim \sum_{l=1}^{N_L-1} \tilde{\beta}_l(A_t) \delta_{l/N_L}(\cdot).$$

- If $R_t^m = 1$, compute the parameters $\hat{\lambda}_j(A_t)$ for $j = 1, \dots, d$ and sample

$$D_t^m \sim \text{Mult}(\hat{\lambda}_1(A_t), \dots, \hat{\lambda}_d(A_t)),$$

and set $L_t^m = 1/2$.

Divide the current node A_t with $J_t^m = (D_t^m, L_t^m)$ to update the tree T_t^m .

6. Storing the information of the state's posterior

Given J_t^m , compute $\varphi_i(A_t)$ for $i = 1, \dots, I$ and store them to the memory.

7. Updating the importance weight

Compute the incremental weight $w_t(T_{t-1}^m, R_t^{m,p})$ and update the importance weights as

$$W_t^m = \frac{W_{t-1}^m w_t(T_{t-1}^m, R_t^{m,p})}{\sum_{m'=1}^M W_{t-1}^{m'} w_t(T_{t-1}^{m'}, R_t^{m',p})}.$$

If the effective sample size $1 / \sum_{m=1}^M (W_t^m)^2$ is less than some prespecified threshold (M/10, say), resample the particles.

B Additional algorithm for the MRS model

To describe the algorithm proposed in Soriano and Ma (2017), we keep using the same notations in Section 3.2. Given the tree structure T , we compute functions $\tilde{\psi}(A)$ for

$A \in \mathcal{N}(T)$ in the bottom-up (from the leaf nodes to the root node) manner as follows:

$$\tilde{\psi}(A) = \begin{cases} \tilde{\xi}_{2,2}(A) + \tilde{\xi}_{2,3}(A) & \text{if } A \in \mathcal{L}(T), \\ \tilde{\xi}_{2,2}(A)\tilde{\psi}(A_l)\tilde{\psi}(A_r) + \tilde{\xi}_{2,3}(A) & \text{if } A \in \mathcal{N}(T) \setminus \{\Omega\}, \\ \tilde{\xi}_{1,2}(A)\tilde{\psi}(A_l)\tilde{\psi}(A_r) + \tilde{\xi}_{1,3}(A) & \text{if } A = \Omega. \end{cases}$$

Recall that only the first row of $\xi(\Omega)$ is meaningful as the initial distribution. Then we obtain $\tilde{\psi}(\Omega) = P(H_0 | T, \mathbf{x})$.

C Proofs

C.1 Posterior computation

Lemma C.1. *For the finite tree T_t , let A_s be a node whose children nodes are leaf nodes. Then we have*

$$\pi_t(V_s = i | T_t) = \frac{\pi_t(T_t, V_s = i)}{\pi_t(T_t)} = \varphi_i(A_s),$$

where $\varphi_i(A_s)$ is defined in (5).

(Proof) Suppose that A_s belongs to the k th layer of T_t . Then there is a sub-sequence $\{\rho(l)\}_{l=1}^k$ such that $A_{\rho(l)}$ belongs to the l th layer and

$$\Omega = A_{\rho(1)} \supset A_{\rho(2)} \supset \cdots \supset A_{\rho(k)} = A_s.$$

By the definition of $\pi_t(T_t, V_{1:t})$, for a sequence $\{v_l\}_{l=1}^k$ such that $v_l \in \{1, \dots, I\}$, we

obtain the expression of the conditional posterior of $\{V_{\rho(l)}\}_{l=1}^k$ as

$$\begin{aligned}\pi_t(\{V_{\rho(l)}\}_{l=1}^k = \{v_l\}_{l=1}^k \mid T_t) &\propto P(\{V_{\rho(l)}\}_{l=1}^k = \{v_l\}_{l=1}^k) \prod_{l=1}^k M_{v_l}(A_{\rho(l)} \mid J_{\rho(l)}) \\ &= \prod_{l=1}^k \boldsymbol{\xi}_{v_{l-1}, v_l}(A_{\rho(l)}) M_{v_l}(A_{\rho(l)} \mid J_{\rho(l)}),\end{aligned}$$

where $v_0 = 1$. We show that for every $k = 1, \dots, K$

$$\begin{aligned}\pi_t(V_{\rho(k)} = v_k \mid T_t) &\propto \sum_{v_1=1}^I \cdots \sum_{v_{k-1}=1}^I \left\{ \prod_{l=1}^k \boldsymbol{\xi}_{v_{l-1}, v_l}(A_{\rho(l)}) M_{v_l}(A_{\rho(l)} \mid J_{\rho(l)}) \right\} \\ &\propto \boldsymbol{\varphi}_{v_k}(A_s)\end{aligned}\tag{13}$$

holds by induction. First, if $k = 1$, which is equivalent to $s = 1$, $\rho(1) = 1$, and $A_s = \Omega$, the posterior of $V(\Omega)$ is written as

$$\pi_t(V(\Omega) = v_1 \mid T_t) \propto \boldsymbol{\xi}_{1, v_1}(\Omega) M_{v_1}(A_1 \mid J_1) \propto \boldsymbol{\varphi}_{v_1}(\Omega).$$

Second, assume that (13) holds for $k = \bar{k}$. Then, if $k = \bar{k} + 1$, we have

$$\begin{aligned}\pi_t(V_{\bar{k}} = v_{\bar{k}}, V_{\bar{k}+1} = v_{\bar{k}+1} \mid T_t) &\propto \sum_{v_1=1}^I \cdots \sum_{v_{\bar{k}-1}=1}^I \pi_t(\{V_{\rho(l)}\}_{l=1}^{\bar{k}+1} = \{v_l\}_{l=1}^{\bar{k}+1} \mid T_t) \\ &\propto \sum_{v_1=1}^I \cdots \sum_{v_{\bar{k}-1}=1}^I \left\{ \prod_{l=1}^{\bar{k}} \boldsymbol{\xi}_{v_{l-1}, v_l}(A_{\rho(l)}) M_{v_l}(A_{\rho(l)} \mid J_{\rho(l)}) \right\} \boldsymbol{\xi}_{v_{\bar{k}}, v_{\bar{k}+1}}(A_s) M_{v_{\bar{k}+1}}(A_s \mid J_s) \\ &\propto \boldsymbol{\varphi}_{v_{\bar{k}}}(A_{\rho(\bar{k})}) \boldsymbol{\xi}_{v_{\bar{k}}, v_{\bar{k}+1}}(A_s) M_{v_{\bar{k}+1}}(A_s \mid J_s),\end{aligned}$$

from which we obtain

$$\begin{aligned}
\pi_t(V_{\bar{k}+1} = v_{\bar{k}+1} \mid T_t) &= \sum_{v_{\bar{k}}=1}^I \pi_t(V_{\bar{k}} = v_{\bar{k}}, V_{\bar{k}+1} = v_{\bar{k}+1} \mid T_t) \\
&\propto \sum_{v_{\bar{k}}=1}^I \varphi_{v_{\bar{k}}}(A_{\rho(\bar{k})}) \xi_{v_{\bar{k}}, v_{\bar{k}+1}}(A_s) M_{v_{\bar{k}+1}}(A_s \mid J_s) \\
&\propto \varphi_{v_{\bar{k}+1}}(A_s). \quad \square
\end{aligned}$$

Proof of Proposition 3.1

Let the finite tree T_t consist of a sequence of decisions $J_{1:t} = \{J_s\}_{s=1}^t$, which sequentially divides nodes $A_{1:t} = \{A_s\}_{s=1}^t$. To derive the proposition for the marginal posterior of T_t , we first consider the joint posterior of T_t and a sequence of the state variables $V_{1:t}$ which are defined for the nodes $A_{1:t}$. From the structure of the model, the joint posterior is written as

$$\begin{aligned}
\pi_t(T_t, V_{1:t}) &= P(J_{1:t}, V_{1:t} \mid \mathbf{x}) \\
&= \frac{1}{Z_t} P(J_{1:t}) P(V_{1:t}) \prod_{s=1, A_s \in \mathcal{N}(T_t)}^t M_{V_s}(A_s \mid J_s) \prod_{s=1, A_s \in \mathcal{L}(T_t)}^t \mu(\mathbf{x}(A_s) \mid A_s),
\end{aligned} \tag{14}$$

where Z_t is the normalizing constant, and $A_{s,l}$ and $A_{s,r}$ are the children nodes of A_s .

For T_{t-1} and T_t , since A_t is divided into $A_{t,l}$ and $A_{t,r}$, we have

$$\begin{aligned}
\mathcal{N}(T_t) &= \mathcal{N}(T_{t-1}) \cup \{A_t\}, \\
\mathcal{L}(T_t) &= \mathcal{L}(T_{t-1}) \setminus \{A_t\} \cup \{A_{t,l}, A_{t,r}\}.
\end{aligned}$$

With the expression of the joint posterior in (14), we obtain

$$\pi_t(T_t, V_{1:t}) = \frac{Z_t}{Z_{t-1}} \pi_{t-1}(T_{t-1}, V_{1:t-1}) P(J_t) P(V_t | V_{1:t-1}) M_{V_t}(A_t | J_t) \frac{\mu(\mathbf{x}(A_{s,l}) | A_{t,l}) \mu(\mathbf{x}(A_{s,r}) | A_{s,r})}{\mu(\mathbf{x}(A_s) | A_s)}. \quad (15)$$

Let A_t^p denote the parent node of A_t and $V_t^p = V(A_t^p)$. Then, since the state variables follow the hidden Markov process, $P(V_t | V_{1:t-1}) = \boldsymbol{\xi}_{V_t^p, V_t}(A_t)$. Integrating out $V_{1:t-1} \setminus V_t^p$ in (15) gives

$$\pi_t(T_t, V_t^p, V_t) = \frac{Z_t}{Z_{t-1}} \pi_{t-1}(T_{t-1}, V_t^p) P(J_t) \boldsymbol{\xi}_{V_t^p, V_t}(A_t) M_{V_t}(A_t | J_t) \frac{\mu(\mathbf{x}(A_{s,l}) | A_{t,l}) \mu(\mathbf{x}(A_{s,r}) | A_{s,r})}{\mu(\mathbf{x}(A_s) | A_s)}.$$

Because A_t is a leaf node of T_{t-1} , by Lemma C.1, we have

$$\begin{aligned} \pi_{t-1}(T_{t-1}, V_t^p = j) &= \pi_{t-1}(T_{t-1}) \pi_{t-1}(V_t^p = j | T_{t-1}) \\ &= \pi_{t-1}(T_{t-1}) \varphi_j(A_t^p). \end{aligned}$$

Hence, we obtain the expression of the marginal distribution of T_t as

$$\begin{aligned} \pi_t(T_t) &= \sum_{i=1}^I \sum_{j=1}^I \pi_t(T_t, V_t^p = j, V_t = i) \\ &= \frac{Z_t}{Z_{t-1}} \pi_{t-1}(T_{t-1}) P(J_t) \sum_{i=1}^I \left\{ \sum_{j=1}^I \varphi_j(A_t^p) \boldsymbol{\xi}_{j,i}(A_t) \right\} M_i(A_t | J_t) \frac{\mu(\mathbf{x}(A_{t,l}) | A_{t,l}) \mu(\mathbf{x}(A_{t,r}) | A_{t,r})}{\mu(\mathbf{x}(A_t) | A_t)}, \end{aligned}$$

which completes the proof. □

Proof of Corollary 3.1

For $\pi_t(D_t | T_{t-1})$, by Proposition 3.1, we obtain

$$\begin{aligned} \pi_t(D_t = j | T_{t-1}) &= \sum_{l=1}^{N_L-1} \pi_t((j, l/N_L) | T_{t-1}) \\ &\propto \sum_{l=1}^{N_L-1} P(D_t = j, L_t = l/N_L) h((j, l/N_L) | A_t) \\ &\propto \lambda_j(A_t) \sum_{l=1}^{N_L-1} \beta_l(A_t) h((j, l/N_L) | A_t). \end{aligned}$$

On the other hand, the conditional probability of L_t is obtained as follows:

$$\begin{aligned} \pi_t(L_t = l/N_L | D_t = j, T_{t-1}) &\propto P(D_t = j, L_t = l/N_L) h((j, l/N_L) | A_t) \\ &\propto \beta(A_t) h(j, l/N_L | T_{t-1}). \end{aligned}$$

The expression of $w_t(T_{t-1})$ immediately follows Proposition 3.1. □

Proof of Proposition 3.2

In this discussion, we suppress \mathbf{x} and T^m in the expectation for simplicity. First, when $A = \Omega$, by the definition $e_\Omega(i) = \tilde{\gamma}_{1,i}(\Omega)$. Next, if A is not the root node, we can decompose $e_A(i')$ as

$$e_A(i') = \sum_{i=1}^I \mathbb{E}[Q(A) I[V(A) = i'] I[V(A^p) = i]].$$

For the summand, because $\theta(A^p)$ and $V(A)$ are conditionally independent given $V(A^p)$, we obtain

$$\begin{aligned}
& \mathbb{E}[Q(A)I[V(A) = i']I[V(A^p) = i]] \\
&= \mathbb{E}[\mathbb{E}[\vartheta(A^p)I[V(A) = i'] \mid V(A^p)]I[V(A^p) = i]Q(A^p)] \\
&= \mathbb{E}[\mathbb{E}[\vartheta(A^p)I[V(A) = i'] \mid V(A^p) = i]I[V(A^p) = i]Q(A^p)] \\
&= \mathbb{E}[\mathbb{E}[\vartheta(A^p) \mid V(A^p) = i]P(V(A) = i' \mid V(A^p) = i)I[V(A^p) = i]Q(A^p)] \\
&= \tilde{\xi}_{i,i'}\mathbb{E}[\vartheta(A^p) \mid V(A^p) = i]e_{A^p}(i).
\end{aligned}$$

Therefore, we obtain

$$e_A(i') = \sum_{i=1}^I \tilde{\xi}_{i,i'}(A)\mathbb{E}[\vartheta(A^p) \mid V(A^p) = i]e_{A^p}(i).$$

C.2 Asymptotic properties of the tree posteriors

Lemma C.2. *Let $\{X_n^i\}_{n=1,2,\dots}$ ($i = 1, \dots, L$) be sequences of random variables that satisfy the following conditions:*

1. $X_n^i > 0$ for every i and n .
2. As $n \rightarrow \infty$, the following convergence occurs

$$\frac{\log X_n^i}{n} \xrightarrow{p} c_i,$$

where $c_1 \geq c_2 \geq \dots \geq c_L$.

Then we have

$$\frac{\log \sum_{i=1}^L X_n^i}{n} \xrightarrow{p} c_1.$$

(Proof) We only discuss the case where there exists l such that $c_l > c_{l+1}$. Let $\epsilon > 0$.

Then, by the first condition,

$$P\left(\frac{\log \sum_{i=1}^L X_n^i}{n} < c_1 - \epsilon\right) \leq P\left(\frac{\log X_n^1}{n} < c_1 - \epsilon\right) \rightarrow 0$$

as $n \rightarrow \infty$. On the other hand,

$$\begin{aligned} & P\left(\frac{\log \sum_{i=1}^L X_n^i}{n} < c_1 + \epsilon\right) \\ & \geq P\left(\left\{\frac{\log L \max_{i=1, \dots, L} X_n^i}{n} < c_1 + \epsilon\right\} \cap \bigcap_{i=1}^l \bigcap_{j=l+1}^L \{X_n^i > X_n^j\}\right) \\ & \geq P\left(\left\{\frac{\log L \max_{i=1, \dots, l} X_n^i}{n} < c_1 + \epsilon\right\} \cap \bigcap_{i=1}^l \bigcap_{j=l+1}^L \{X_n^i > X_n^j\}\right) \rightarrow 1. \quad \square \end{aligned}$$

Lemma C.3. For $T \in \mathcal{T}^K$ and $A \in \mathcal{N}(T)$, if $i \in \tau(A | T)$, then

$$\frac{\log M_i(A | j_A)}{n} \xrightarrow{p} \sum_{g=1}^G \zeta_g P_g(A) [P_g(A_l | A) \log P_g(A_l | A) + P_g(A_r | A) \log P_g(A_r | A)],$$

where j_A is the splitting rule that divides A into A_l and A_r .

(Proof) By the result of Schwarz (1978), since the parameter $\theta(A)$ follow the beta distribution, which belongs to a continuous exponential family, the log of the marginal likelihood is written as

$$\begin{aligned} \log M_i(A | T) &= \hat{l}_A(i, T) - \frac{r_i}{2} \log n(A) + \mathcal{O}_p(1), \\ \hat{l}_A(i, T) &= \log \left[\prod_{g=1}^G \hat{\theta}_g(A)^{n_g(A_l)} (1 - \hat{\theta}_g(A))^{n_g(A_r)} \right] \\ &= \sum_{g=1}^G \left[n_g(A_l) \log \hat{\theta}_g(A) + n_g(A_r) \log(1 - \hat{\theta}_g(A)) \right], \end{aligned} \quad (16)$$

where the definition of $\hat{\theta}_g(A)$ (the MLE) and r_i (the number of parameters) depend on which type of priors in Assumption 1 is introduced by the state i :

$$\hat{\theta}_g(A) = \begin{cases} \frac{n_g(A_l)}{n_g(A)} & \text{(Prior A),} \\ \frac{n(A_l)}{n(A)} & \text{(Prior B),} \\ \frac{P_g(A_l)}{P_g(A)} & \text{(Prior C),} \end{cases} \quad r_i = \begin{cases} G & \text{(Prior A),} \\ 1 & \text{(Prior B),} \\ 0 & \text{(Prior C).} \end{cases} \quad (17)$$

Notice that, for Prior C, the constant $c(A)$ and the true measures P_g must satisfy

$$c(A) = P_g(A_l | A)$$

for every $g = 1, \dots, G$ because if this does not hold, the state i is not included in the set of feasible states $\tau(A | T)$.

Since $i \in \tau(A | T)$, the law of large numbers gives $\hat{\theta}_g(A) \xrightarrow{p} P_g(A_l | A)$. Hence, we obtain the limit of $\hat{l}_A(i, T)/n$ as

$$\begin{aligned} \frac{\hat{l}_A(i, T)}{n} &= \sum_{g=1}^G \frac{n_g(\Omega)}{n(\Omega)} \frac{n_g(A)}{n_g(\Omega)} \left[\frac{n_g(A_l)}{n_g(A)} \log \hat{\theta}_g(A) + \frac{n_g(A_r)}{n_g(A)} \log(1 - \hat{\theta}_g(A)) \right] \\ &\xrightarrow{p} \sum_{g=1}^G \zeta_g P_g(A) [P_g(A_l | A) \log P_g(A_l | A) + P_g(A_r | A) \log P_g(A_r | A)]. \end{aligned}$$

□

Lemma C.4. For $T \in \mathcal{T}^K$, $A \in \mathcal{N}(T)$, $i \in \tau(A | T)$ and $j \in \{1, \dots, I\}$, we have

$$\frac{\log M_i(A | T) - \log M_j(A | T)}{n} \xrightarrow{p} c_{i,j},$$

where $c_{i,j} = 0$ if $j \in \tau(A | T)$ and $c_{i,j} > 0$ if $j \in \{1, \dots, I\} \setminus \tau(A | T)$.

(Proof) If $j \in \tau(A | T)$, obtaining the result

$$\frac{\log M_i(A | T) - \log M_j(A | T)}{n} \xrightarrow{p} 0$$

is straightforward from the proof of Proposition C.3. Hence we consider the case of $j \in \{1, \dots, I\} \setminus \tau(A | T)$. Under the state j , for every g , the estimator $\hat{\theta}_g(A)$ is defined as in (17), and there exists $C_g \in (0, 1)$ such that $\hat{\theta}_g(A) \xrightarrow{p} C_g$. By the definition of $\tau(A | T)$, there exists g^* such that $C_{g^*} \neq P_{g^*}(A_l | A)$. As in the proof of Proposition C.3, for the difference of the marginal likelihoods, we obtain

$$\begin{aligned} \frac{\log M_i(A | T) - \log M_j(A | T)}{n} &\xrightarrow{p} \sum_{g=1}^G \zeta_g P_g(A) \Lambda_g, \\ \Lambda_g &= P_g(A_l | A) \log \frac{P_g(A_l | A)}{C_g} + P_g(A_r | A) \log \frac{P_g(A_r | A)}{1 - C_g}. \end{aligned}$$

Because Λ_g is the KL divergence of the two discrete distributions, $\Lambda_g \geq 0$ for all g and $\Lambda_{g^*} > 0$. By Assumption 1, this result implies that

$$\sum_{g=1}^G \zeta_g P_g(A) \Lambda_g > 0.$$

□

Proof of Theorem 4.1 (non-informative prior) and Theorem 4.2

This section provides the proofs of Theorem 4.1 under the assumption that the prior of the location variables is independent of the sample size, that is, $\eta = 0$ in Eq (2). The general case is discussed in the next section.

In this proof, we modify the notation for the marginal likelihood defined in Eq. (4) and use $M_i(A | T)$ to represent the likelihood on A of the tree T under the i th state

to reflect its dependency on the tree structure.

Let $T \in \mathcal{T}^K$ and \mathcal{V} denote a set of a combination of the states for all of the non-leaf nodes of T . Notice that an element of \mathbf{V} does not need to satisfy $P(\mathbf{V} = \mathbf{v}) > 0$, where \mathbf{V} is the totality of the state variables. In the following proof, for $\mathbf{v} \in \mathcal{V}$, $\mathbf{v}(A)$ denotes a state on a node A . Let $l(\mathbf{v}, T)$ denote the log of the joint likelihood function

$$l(\mathbf{v}, T) = \log P(\mathbf{x}_n | T, \mathbf{v}) = \sum_{A \in \mathcal{N}(T)} l_A(\mathbf{v}(A), T) + \sum_{A \in \mathcal{L}(T)} \log \mu(\mathbf{x}_n(A) | A), \quad (18)$$

where $l_A(\mathbf{v}(A), T) = \log M_A(\mathbf{v}(A) | T)$. By Schwarz (1978), this likelihood l_A has the following expression

$$l_A(\mathbf{v}(A), T) = \hat{l}_A(\mathbf{v}(A), T) - \frac{r_{\mathbf{v}(A)}}{2} \log n(A) + \mathcal{O}_p(1),$$

where \hat{l}_A and r_i are defined in (16). Let $\bar{\mathbf{v}} \in \mathcal{V}$ be a collection of states such that, for all $A \in \mathcal{N}(T)$, $\theta_g(A)$ is fixed to $\mu(A_l)/\mu(A)$. For $\bar{\mathbf{v}}$, we have

$$\begin{aligned} l_A(\bar{\mathbf{v}}, T) &= \sum_{A \in \mathcal{N}(T)} \left\{ n(A_l) \log \left(\frac{\mu(A_l)}{\mu(A)} \right) + n(A_r) \log \left(\frac{\mu(A_r)}{\mu(A)} \right) \right\} + \sum_{A \in \mathcal{L}(T)} \log \mu(\mathbf{x}_n(A) | A) \\ &= \log \mu(\mathbf{x}) = 0. \end{aligned}$$

Hence $l(\mathbf{v}, T)$ is rewritten as

$$l(\mathbf{v}, T) = l(\mathbf{v}, T) - l(\bar{\mathbf{v}}, T) = \sum_{A \in \mathcal{N}(T)} \left\{ \hat{l}_A(\mathbf{v}(A), T) - \hat{l}_A(\bar{\mathbf{v}}(A), T) \right\} - \frac{C(\mathbf{v})}{2} \log n + \mathcal{O}_p(1).$$

For the part inside of the braces, when \mathbf{v} is replaced with $\mathbf{v}_T \in \mathcal{V}_T$, where $\mathcal{V}_T = \{\mathbf{v} :$

$\mathbf{v}(A) \in \tau(A | T)$ for all $A \in \mathcal{N}(T)$, the definition of \hat{l}_A gives

$$\begin{aligned} \frac{\hat{l}(\mathbf{v}_T(A), T) - \hat{l}(\bar{\mathbf{v}}(A), T)}{n} &= \sum_{g=1}^G \frac{n_g(A)}{n} \left[\frac{n_g(A_l)}{n_g(A)} \log \frac{\hat{\theta}_g(A)}{\mu(A_l | A)} + \frac{n_g(A_r)}{n_g(A)} \log \frac{1 - \hat{\theta}_g(A)}{\mu(A_r | A)} \right] \\ &\xrightarrow{p} \sum_{g=1}^G \zeta_g P_g(A) \left[P_g(A_l | A) \log \frac{P_g(A_l | A)}{\mu(A_l | A)} + P_g(A_r | A) \log \frac{P_g(A_r | A)}{\mu(A_r | A)} \right]. \end{aligned}$$

For all $A \in \mathcal{L}(T)$, there exists an unique sequence of nodes

$$\Omega = B_{A,0} \supset B_{A,1} \supset \cdots \supset B_{A,K} = A, \quad (19)$$

where $B_{A,k} \in T$ ($k = 0, \dots, K$) is a node in the k th level. With this sequence, we obtain the limit of the scaled log-likelihood as

$$\begin{aligned} &\frac{l(\mathbf{v}_T, T)}{n} \\ &\xrightarrow{p} \sum_{A \in \mathcal{N}(T)} \sum_{g=1}^G \zeta_g P_g(A) \left[P_g(A_l | A) \log \frac{P_g(A_l | A)}{\mu(A_l | A)} + P_g(A_r | A) \log \frac{P_g(A_r | A)}{\mu(A_r | A)} \right] \\ &= \sum_{g=1}^G \zeta_g \sum_{A \in \mathcal{N}(T)} \left[P_g(A_l) \log \frac{P_g(A_l | A)}{\mu(A_l | A)} + P_g(A_r) \log \frac{P_g(A_r | A)}{\mu(A_r | A)} \right] \\ &= \sum_{g=1}^G \zeta_g \sum_{A \in \mathcal{L}(T)} P_g(A) \left[\log \frac{P_g(B_{A,1} | B_{A,0})}{\mu(B_{A,1} | B_{A,0})} + \cdots + \log \frac{P_g(B_{A,K} | B_{A,K-1})}{\mu(B_{A,K} | B_{A,K-1})} \right] \\ &= \sum_{g=1}^G \zeta_g \sum_{A \in \mathcal{L}(T)} P_g(A) \log \frac{P_g(A)}{\mu(A)} = \sum_{g=1}^G \zeta_g KL(P_g|_T || \mu). \end{aligned} \quad (20)$$

Because $P_g|_T$ admits the density function

$$p_g|_T(x) = \sum_{A \in \mathcal{L}(T)} \mathbf{1}_A(x) \frac{P_g(A)}{\mu(A)}, \quad x \in \Omega$$

the KL divergence in (20) is rewritten as

$$\begin{aligned}
KL(P_g|_T||\mu) &= \sum_{A \in \mathcal{L}(T)} P_g(A) \log \frac{P_g(A)}{\mu(A)} \\
&= \int p_g \sum_{A \in \mathcal{L}(T)} \mathbf{1}_A \log \frac{P_g(A)}{\mu(A)} d\mu \\
&= \int p_g \log p_g|_T d\mu \\
&= \int p_g(x) \log \frac{p_g(x)}{\mu(x)} d\mu(x) - \int p_g \log \frac{p_g}{p_g|_T} d\mu \\
&= KL(P_g||\mu) - KL(P_g||P_g|_T).
\end{aligned}$$

Because $KL(P_g||\mu)$ is independent of T , we obtain another expression of \mathcal{T}_M^K in (10) as

$$\mathcal{T}_M^K = \arg \max_{T \in \mathcal{T}^K} \sum_{g=1}^G \zeta_g KL(P_g|_T||\mu)$$

By Lemma C.4 and (18), for $\mathbf{v} \in \mathcal{V} \setminus \mathcal{V}_T$, we can show that

$$\text{p-lim}_{n \rightarrow \infty} \frac{l(\mathbf{v}, T) - l(\mathbf{v}_T, T)}{n} = \text{p-lim}_{n \rightarrow \infty} \sum_{A \in \mathcal{N}(T)} \frac{l_A(\mathbf{v}, T) - l_A(\mathbf{v}_T, T)}{n} > 0, \quad (21)$$

and $\text{p-lim}_{n \rightarrow \infty} l(\mathbf{v}, T)/n$ exists. Hence, for $T_M \in \mathcal{T}_M^K$, $\mathbf{v}' \in \mathcal{V}_{T_M}$, $T \in \mathcal{T}^K \setminus \mathcal{T}_M^K$ and $\mathbf{v} \in \mathcal{V}$, we have

$$\text{p-lim}_{n \rightarrow \infty} \frac{l(\mathbf{v}', T_M) - l(\mathbf{v}, T)}{n} \geq \sum_{g=1}^G \zeta_g KL(P_g|_{T_M}||\mu) - \sum_{g=1}^G \zeta_g KL(P_g|_T||\mu) > 0.$$

Hence, for such T_M and \mathbf{v}' , we obtain

$$\frac{P(\mathbf{x}_n | T)}{P(\mathbf{x}_n | T_M)} = \frac{P(T) \sum_{\mathbf{v} \in \mathcal{V}} \exp(l(\mathbf{v}, T)) P(\mathbf{v})}{P(T_M) \sum_{\mathbf{v} \in \mathcal{V}} \exp(l(\mathbf{v}, T_M)) P(\mathbf{v})} \leq \sum_{\mathbf{v} \in \mathcal{V}} \frac{\exp(l(\mathbf{v}, T)) P(T) P(\mathbf{v})}{\exp(l(\mathbf{v}', T_M)) P(T_M) P(\mathbf{v}')} \xrightarrow{p} 0.$$

This result implies $p(T \in \mathcal{T}_M^K \mid \mathbf{x}_n) \xrightarrow{p} 1$, which completes the proof of Theorem 4.1.

To prove Theorem 4.2, we fix $T \in \mathcal{T}^K$ and define a set \mathcal{S}_T as

$$\mathcal{S}_T = \left\{ \mathbf{v} \in \mathcal{V}_T \mid \mathbf{v} \in \arg \min_{\mathbf{v}' \in \mathcal{V}_T} C(\mathbf{v}') \right\}.$$

Then we want to show $P(\mathbf{V} \in \mathcal{S}_T \mid T, \mathbf{x}_n) \xrightarrow{p} 1$. The result (21) implies

$$p(\mathbf{V} \in \mathcal{V}_T \mid T, \mathbf{x}_n) \xrightarrow{p} 1,$$

so we only need to compare the elements of \mathcal{V}_T . Let $\mathbf{v} \in \mathcal{V}_T \setminus \mathcal{S}_T$ and $\mathbf{v}' \in \mathcal{S}_T$. For the difference of the log likelihoods, we have

$$l(\mathbf{v}', T) - l(\mathbf{v}, T) = \sum_{A \in \mathcal{N}(T)} \left[\hat{l}_A(\mathbf{v}'(A), T) - \hat{l}_A(\mathbf{v}(A), T) \right] + \frac{C(\mathbf{v}) - C(\mathbf{v}')}{2} \log n + \mathcal{O}_p(1),$$

where \hat{l}_A and C is defined in (17) and (11), respectively. If $\mathbf{v}(A)$ and $\mathbf{v}'(A)$ introduce the same type of the prior (e.g., Prior A and Prior A), because the corresponding estimators $\hat{\theta}_g(A)$ have the same form,

$$\hat{l}_A(\mathbf{v}'(A), T) - \hat{l}_A(\mathbf{v}(A), T) = 0.$$

On the other hand, if $\mathbf{v}(A)$ and $\mathbf{v}'(A)$ introduce different types of the prior (e.g., Prior A and Prior B), because they are the maximized log-likelihood under the two nested hypotheses,

$$-2[\hat{l}_A(\mathbf{v}'(A), T) - \hat{l}_A(\mathbf{v}(A), T)]$$

weakly converges to the χ^2 distribution (Wilks, 1938). Hence, we obtain

$$\frac{l(\mathbf{v}', T) - l(\mathbf{v}, T)}{\log n} \xrightarrow{p} \frac{C(\mathbf{v}) - C(\mathbf{v}')}{2} > 0,$$

which implies $P(\mathbf{V} \in \mathcal{S}_T \mid T, \mathbf{x}_n) \xrightarrow{p} 1$. □

Proof of Theorem 4.1 (informative prior)

This section provides the proofs of Theorem 4.1 for the case of the informative prior ($\eta > 0$). In the proof, we often use the results provided in the previous section on the non-informative prior and use the same notations.

We first discuss the asymptotic behavior of the prior of trees, which is in this case dependent on the sample size n . We only need to focus on the prior of the location variables L since this is only the component that depends on the sample size as given in Eq (2). By the fact that

$$\frac{n(A)}{n} \xrightarrow{p} \sum_{g=1}^G \zeta_g P_g(A) \tag{22}$$

and Lemma C.2, we obtain the limit

$$\frac{\log P(L(A) = \mu(A_l)/\mu(A))}{n} \xrightarrow{p} -\eta \sum_{g=1}^G \zeta_g P_g(A) f\left(\left|\frac{\mu(A_l)}{\mu(A)} - 0.5\right|\right).$$

Since given $T \in \mathcal{T}^K$, the location variables are all independent and the prior of the dimension variables is independent of the sample size n , the log-tree prior $l(T) = \log P(T)$ has a limit as follows:

$$\frac{l(T)}{n} \xrightarrow{p} -\eta \sum_{g=1}^G \zeta_g B_g(T),$$

which is the penalty term introduced in Theorem 4.1.

We next define $\psi(T)$ as

$$\psi(T) = \sum_{g=1}^G \zeta_g \{KL(P_g|_T||\mu) - \eta B_g(T)\}.$$

By the the proof of Theorem 4.1 for the non-informative case, this is a limit under $\mathbf{v} \in \mathcal{V}_T$ (this \mathcal{V}_T , and \mathcal{V} , a collection of possible combination of the state variables on the tree, are defined in the proof for the non-informative case), that is,

$$\frac{l(T) + l(\mathbf{v}, T)}{n} \xrightarrow{p} \psi(T)$$

as $n \rightarrow \infty$. On the other hand, for $\mathbf{v} \in \mathcal{V} \setminus \mathcal{V}_T$, the limit exists and

$$\text{p-lim} \frac{l(T) + l(\mathbf{v}, T)}{n} < \psi(T).$$

As discussed in the proof of Theorem 4.1, we have the equivalence

$$\arg \max_{T \in \mathcal{T}^K} \sum_{g=1}^G \zeta_g KL(P_g|_T||\mu) = \arg \min_{T \in \mathcal{T}^K} \sum_{g=1}^G \zeta_g KL(P_g||P_g|_T),$$

which implies that

$$T_M^K = \arg \max_{T \in \mathcal{T}^K} \psi(T).$$

Hence, for $T_M \in \mathcal{T}_M^K$, $\mathbf{v}' \in \mathcal{V}_{T_M}$, $T \in \mathcal{T}^K \setminus \mathcal{T}_M^K$, and $\mathbf{v} \in \mathcal{V}$, we have

$$\text{p-lim}_{n \rightarrow \infty} \frac{(l(T_M) + l(\mathbf{v}', T_M)) - (l(T) + l(\mathbf{v}, T))}{n} \geq \psi(T_M) - \psi(T) > 0.$$

Hence, for such T, \mathcal{T}_M and \mathbf{v}' , it follows that

$$\begin{aligned} \frac{P(T, \mathbf{x}_n)}{P(T_M, \mathbf{x}_n)} &= \frac{P(T)P(\mathbf{x}_n | T)}{P(T_M)P(\mathbf{x}_n | T_M)} = \frac{P(T) \sum_{\mathbf{v} \in \mathcal{V}} P(\mathbf{v}) \exp(l(\mathbf{v}, T))}{P(T_M) \sum_{\mathbf{v} \in \mathcal{V}} P(\mathbf{v}) \exp(l(\mathbf{v}, T_M))} \\ &\leq \sum_{\mathbf{v} \in \mathcal{V}} \frac{\exp(l(T) + l(\mathbf{v}, T))P(\mathbf{v})}{\exp(l(T_M) + l(\mathbf{v}', T_M))P(\mathbf{v}')} \xrightarrow{p} 0. \end{aligned}$$

Therefore, $P(T \in \mathcal{T}_M^K | \mathbf{x}_n) \xrightarrow{p} 1$. □

D Consistency for the MRS model

To describe the consistency, for a possible node A , we define a variable $Z(A)$ as follows:

$$Z(A) = \begin{cases} 1 & \text{if } V(A) = 1, \\ 0 & \text{if } V(A) \in \{2, 3\}. \end{cases}$$

Hence, $\theta_1(A) = \theta_2(A)$ if $Z(A) = 0$ and $\theta_1(A) \neq \theta_2(A)$ with probability one if $Z(A) = 1$. Then we can obtain the following consistency result.

Corollary D.1. *Let $\mathbf{Z} = \{Z(A)\}_{A \in \mathcal{N}(T)}$ and $\mathbf{z} = \{z(A)\}_{A \in \mathcal{N}(T)}$ be a collection of $Z(A)$ on $T \in \mathcal{T}^K$ and one of its realizations, respectively. If $P(\mathbf{Z} = \mathbf{z}) > 0$ for any possible \mathbf{z} , then*

$$P(Z(A) = \mathbf{1}_{\{P_1(A_l|A) \neq P_2(A_l|A)\}} \text{ for all } A \in \mathcal{N}(T) | T, \mathbf{x}_n) \xrightarrow{p} 1,$$

where $\mathbf{1}$ is the indicator function, and

$$P(Z(A) = \mathbf{1}_{\{P_1(A_l|A) \neq P_2(A_l|A)\}} \text{ for all } A \in \mathcal{N}(T) | \mathbf{x}_n) \xrightarrow{p} 1,$$

where T is random.

(Proof) In this case, \mathcal{V}_T in Theorem 4.2 is written as

$$\mathcal{V}_T = \{\mathbf{v} \mid \mathbf{v}(A) = 1 \text{ if } P_1(A_l \mid A) \neq P_2(A_l \mid A)\}.$$

We additionally define $\tilde{\mathcal{V}}_T$ as

$$\tilde{\mathcal{V}}_T = \{\mathbf{v} \mid \mathbf{v}(A) = 2 \text{ if } P_1(A_l \mid A) = P_2(A_l \mid A)\}.$$

Then, under the condition that $\mathbf{v} \in \mathcal{V}_T$, the complexity $C(\mathbf{v})$ is minimized if and only if $\mathbf{v} \in \tilde{\mathcal{V}}_T$. Hence, by Theorem 4.2 we obtain

$$\begin{aligned} P(Z(A) = \mathbf{1}_{\{P_1(A_l|A) \neq P_2(A_l|A)\}} \text{ for all } A \in \mathcal{N}(T) \mid T, \mathbf{x}_n) \\ = P(\mathbf{V} \in \mathcal{V}_T \cap \tilde{\mathcal{V}}_T \mid T, \mathbf{x}_n) \xrightarrow{p} 1. \end{aligned}$$

We can show the second result by using Theorem 4.1 as follows:

$$\begin{aligned} P(\mathbf{V} \in \mathcal{V}_T \cap \tilde{\mathcal{V}}_T \mid \mathbf{x}_n) &= \sum_{T \in \mathcal{T}^K} P(\mathbf{V} \in \mathcal{V}_T \cap \tilde{\mathcal{V}}_T \mid T, \mathbf{x}_n) P(T \mid \mathbf{x}_n) \\ &\geq \arg \min_{T_M \in \mathcal{T}_M^K} P(\mathbf{V} \in \mathcal{V}_T \cap \tilde{\mathcal{V}}_T \mid T_M, \mathbf{x}_n) P(T \in \mathcal{T}_M^K \mid \mathbf{x}_n) \\ &\xrightarrow{p} 1. \end{aligned}$$

□

E Numerical evaluation of the density estimators under different complexity of data generating processes

In this section, we observe the behavior of our proposed HMPT algorithm in density estimation in a case where observed points concentrate around certain points or the boundary of the sample space, which occurs in real data analysis in many cases. We consider the following scenario: the variables X_j ($j = 1, \dots, d$) independently follow the beta distribution

$$\text{Beta}(2^a, 2^a),$$

where the parameter a takes values from \mathbb{R} . When $a = 0$, the data is uniformly distributed in the sample space $(0, 1]^d$, and when a is positive, the distribution concentrates on the central point $(0.5, 0.5, \dots, 0.5)$. When a is negative, the distribution concentrates on the vertices of the hyper-cube $(0, 1]^d$. For the cases of $a < 0$, many values generated from the beta distribution tend to be indistinguishable from 0 and 1, and we found it often made the behavior of our HMPT algorithm, which is designed for continuous distributions without mass points, unstable. As such, for the negative a 's, we use an approximative distribution in which the conditional distributions on $(0, 10^{-5}]$ and $(0.9 - 10^{-5}, 1.0]$ are replaced with the uniform.

In the estimation, we used the APT model as described in Section 5.1. The number of particles is 1,000, and the maximum resolution is set to 10 or 15. The data is generated under $a = (-3, -2, -1, 0, 1, 2, 3)$, $n = 10^2, 10^3, 10^4$, and $d = 1, 2, 3, 4$, and for each combination, 50 data sets under different random seeds are generated. The performance is assessed based on L_1 distance between the estimated density (posterior mean) and the true distribution. Since the L_1 distance is difficult to compute

analytically, we use the Monte Carlo approximation with 10000 values generated from the true distributions.

Figure 9 provides the average L_1 distance under the different settings. When the distribution is uniform ($a = 0$), we can see that the L_1 distance is very small, and this can be understood as an advantage of using the APT model to learn the shrinkage level adaptively. On the other hand, especially when $a < 0$ and d is large, the distance is larger for $a \neq 0$ though the performance is slightly improved if the maximum resolution K is larger. One possible reason to explain this result is that the APT model, with which we learn the smoothness of unknown densities, is not designed to capture such extremely spiky distributions. Thus the performance would be improved by modifying the prior distributions but this discussion is out of this research’s scope because our main contribution is proposing the algorithm that works for many types of PT-based models.

Figure 10 compares the number of nodes included in the MAP trees. From this result, we can see that the number is larger for the larger maximum resolution $K = 15$ as naturally expected but tends to be smaller when the parameter a is positive. The latter phenomenon is explained by the rule that the SMC algorithm stops splitting nodes when the number of included observations is below the threshold such as 5. When a is positive, since the distribution concentrates around the central point, nodes with only a few observations tend to be generated in early levels and thus are no longer split, resulting in a smaller number of nodes on the tree. Hence this numerical result clarifies our proposed algorithm’s tendency of “ignoring” nodes with only a few observations “zooming up” nodes including many observations.




Figure 9: The average L_1 distance from the true distribution with error bars that indicate the standard deviation obtained under different sample size, dimensionality, and values of the parameter a .

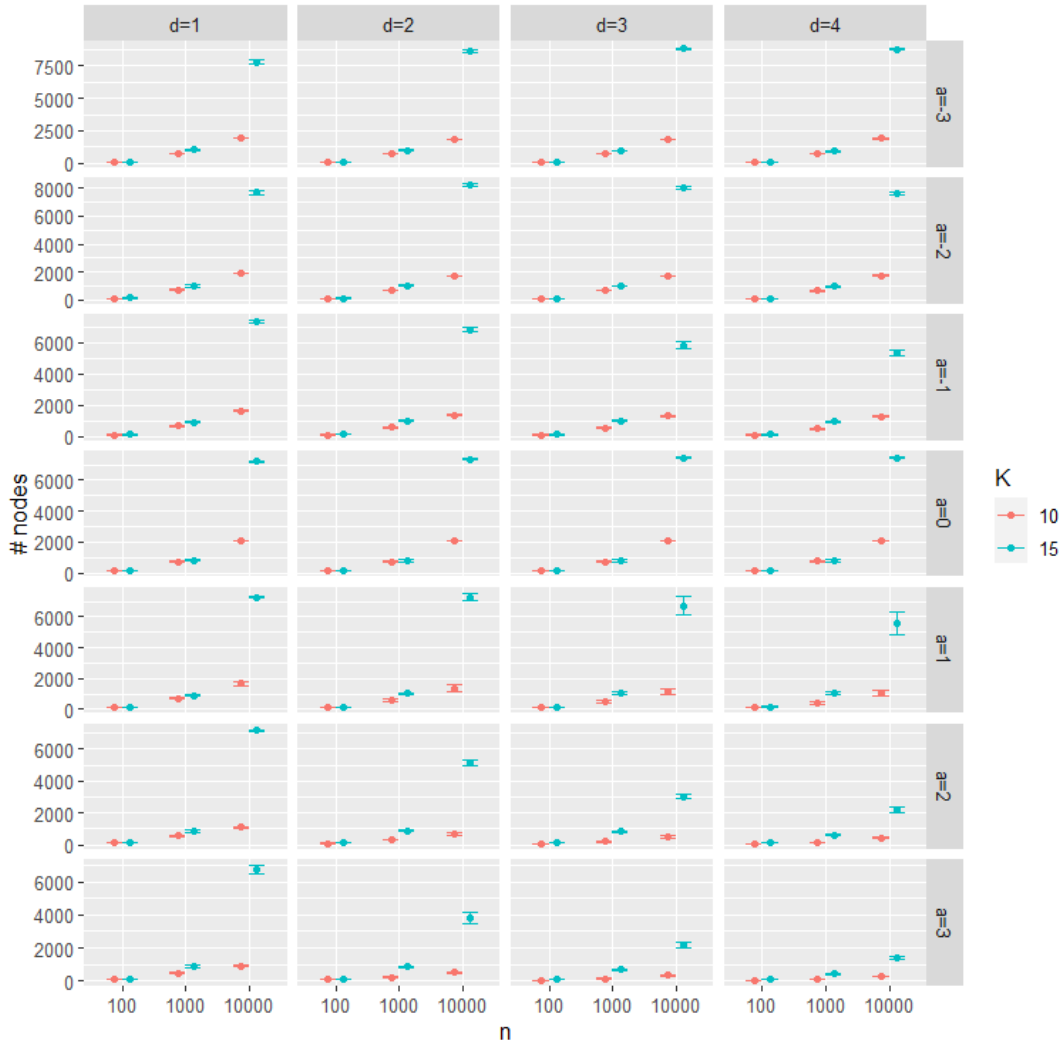


Figure 10: The average number of nodes in the MAP trees with error bars that indicate the standard deviation obtained under different sample size, dimensionality, and values of the parameter a .

F Details of the experiments

F.1 Density estimation

F.1.1 Hyper-parameter settings

For the HMPT model and the original APT model, we use the common settings following Ma (2017). The transition matrix for the latent states on each A that characterize different smoothness levels of the density is given by

$$\xi_{i,i'}(A) \propto \begin{cases} e^{\beta(i-i')} & \text{if } i \leq i', \\ 0 & \text{if } i > i', \end{cases}$$

where $\beta = 0.1$ and the number of states $I = 5$. For the $i (< I)$ th state, given $V(A) = i$, the precision $\nu(A)$ follows the prior

$$\log_{10} \nu(A) \sim \text{Unif}(a(i), a(i+1)],$$

where $a(i) = L + (i-1)(U-L)/(I-1)$ with $L = -1$ and $U = 4$. For the I th state, $F_I = 1_\infty$, so $\nu(A)$ is fixed to ∞ . This is equivalent to stopping the partition and putting the conditional distribution $\mu(\cdot | A)$ on A . In the computation, this uniform distribution is approximated by 5 evenly spaced grid points.

F.1.2 Simulation scenarios in the two-dimensional cases

The data sets are simulated from the following distributions.

1. “Blocks”:
$$\frac{1}{3} \mathbf{1}_{[0.1,0.45] \times [0.35,0.9]}(x_1, x_2) + \frac{1}{3} \mathbf{1}_{[0.2,0.8] \times [0.45,0.5]}(x_1, x_2) + \frac{1}{3} \mathbf{1}_{[0.7,0.9] \times [0.05,0.6]}(x_1, x_2)$$

2. “Clusters”:

$$\begin{aligned} & \frac{1}{10} \text{Beta}(x_1 | 1, 1) \times \text{Beta}(x_2 | 1, 1) + \frac{3}{10} \text{Beta}(x_1 | 15, 45) \times \text{Beta}(x_2 | 15, 45) \\ & + \frac{3}{10} \text{Beta}(x_1 | 45, 15) \times \text{Beta}(x_2 | 22.5, 37.5) + \frac{3}{10} \text{Beta}(x_1 | 37.5, 22.5) \times \text{Beta}(x_2 | 45, 15) \end{aligned}$$

3. “Smooth”: $\text{Beta}(x_1 | 10, 20) \text{Beta}(x_2 | 10, 20)$

F.1.3 Implementation details for the higher dimensional cases

For the algorithm of the original APT (implemented by the `apt` function in the `PTT` package), in the first case with $d = 6$, the maximum resolution is fixed to 9 because setting higher values lead to insufficient memory. Also, because the `apt` function does not scale if the dimension is beyond $d \approx 10$, in the second case with large d , we used the proposed SMC algorithm to carry out inference for the original APT model as well, which corresponds to setting $N_L = 2$.

In the classical PT method, the Dirichlet prior is set to $\text{Dir}(0.1k^2, \dots, 0.1k^2)$ (k : the depth of the node), and the maximum depth is set to 15. For the Dirichlet process Gaussian mixture model, we used the `PYdensity` function in the R package `BNPmix` (Corradin et al., 2021), in which the strength and discounting parameters are set to 10 and 0, respectively, the covariance matrices are all diagonal (`model = "DLS"`), and the size of the burn-in period and the sampling is both set to 1000. For the Gaussian kernel density estimation, it is implemented by the `kde` function in the R package `ks`. We set the bandwidth matrix by using `Hpi.diag` function which selects the optimal diagonal matrix.

F.2 Two-group comparison

F.2.1 Hyper-parameter settings

For the transition matrix $\boldsymbol{\xi}(A)$, we use the form proposed in Soriano and Ma (2017) for incorporating multiple testing control

$$\boldsymbol{\xi}(A) = \begin{bmatrix} (1-\rho)\gamma & (1-\rho)(1-\gamma) & \rho \\ (1-\rho)\gamma 2^{-k} & (1-\rho)(1-\gamma 2^{-k}) & \rho \\ 0 & 0 & 1 \end{bmatrix},$$

where $\gamma \in (0, 1)$, and $\rho \in (0, 1)$, and k is the depth of A , and we set $(\gamma, \rho) = (0.3, 0.3)$ following recommendations in that paper.

F.2.2 Simulation scenarios

1. “Local location shift”: For $j = 1, \dots, 25$,

$$\begin{aligned} (X_{1,2(j-1)+1}, X_{1,2j}) &\sim \frac{1}{3}N(\mu_1, \Sigma) + \sum_{l=2}^3 \frac{1}{3}N(\mu_l, \Sigma), \\ (X_{2,2(j-1)+1}, X_{2,2j}) &\sim \frac{1}{3}N(\mu_1 + \delta_j, \Sigma) + \sum_{l=2}^3 \frac{1}{3}N(\mu_l, \Sigma), \end{aligned}$$

where $\delta_j = -0.5$ for $j = 1, \dots, 5$ and 0 for $j = 6, \dots, 25$.

2. “Local dispersion difference”: For $j = 1, \dots, 25$,

$$\begin{aligned} (X_{1,2(j-1)+1}, X_{1,2j}) &\sim \frac{1}{3}N(\mu_1, \Sigma) + \sum_{l=2}^3 \frac{1}{3}N(\mu_l, \Sigma), \\ (X_{2,2(j-1)+1}, X_{2,2j}) &\sim \frac{1}{3}N(\mu_1, \Sigma + \Delta_j) + \sum_{l=2}^3 \frac{1}{3}N(\mu_l, \Sigma), \end{aligned}$$

where $\Delta_j = -0.4$ for $j = 1, \dots, 5$ and 0 for $j = 6, \dots, 25$.

3. “Correlation”: For $j = 1, \dots, 25$,

$$\begin{aligned} (X_{1,2(j-1)+1}, X_{1,2j}) &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right) \\ (X_{2,2(j-1)+1}, X_{2,2j}) &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \delta_j \\ \delta_j & 1 \end{bmatrix}\right), \end{aligned}$$

where $\delta_j = 0.75$ for $j = 1, \dots, 5$ and $\delta_j = 0$ for $j = 6, \dots, 25$.

In the “local location shift” and “local dispersion difference”, the parameters are

$$\mu_1 = (-2.5, 1.0), \mu_2 = (1.0, -2.0), \mu_3 = (2.0, 2.5), \Sigma = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.7 \end{bmatrix}.$$

G Additional figures

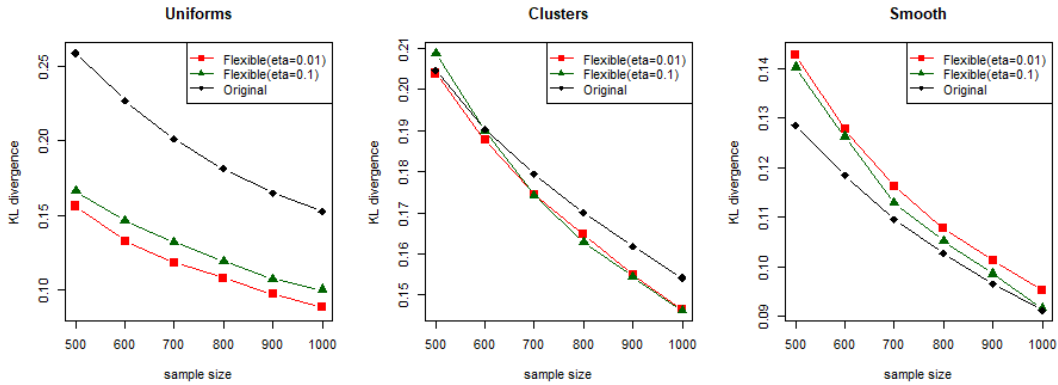


Figure 11: The average KL divergences between the estimated density and the true density.




Figure 12: The comparison of the predictive performance. Each point corresponds to the average of the predictive score in Eq. (12) based on 50 data sets. In the right plot, the predictive scores of the DPM model for the over 30 dimensional cases are below the displayed range.




Figure 13: The MAP tree for the mass cytometry data set. The size and the color indicate the estimated $\text{eff}(A)$. Only the nodes with the sample size larger than 50 are drawn. Since there are a huge number of nodes, the information on the dimension is omitted in this figure.

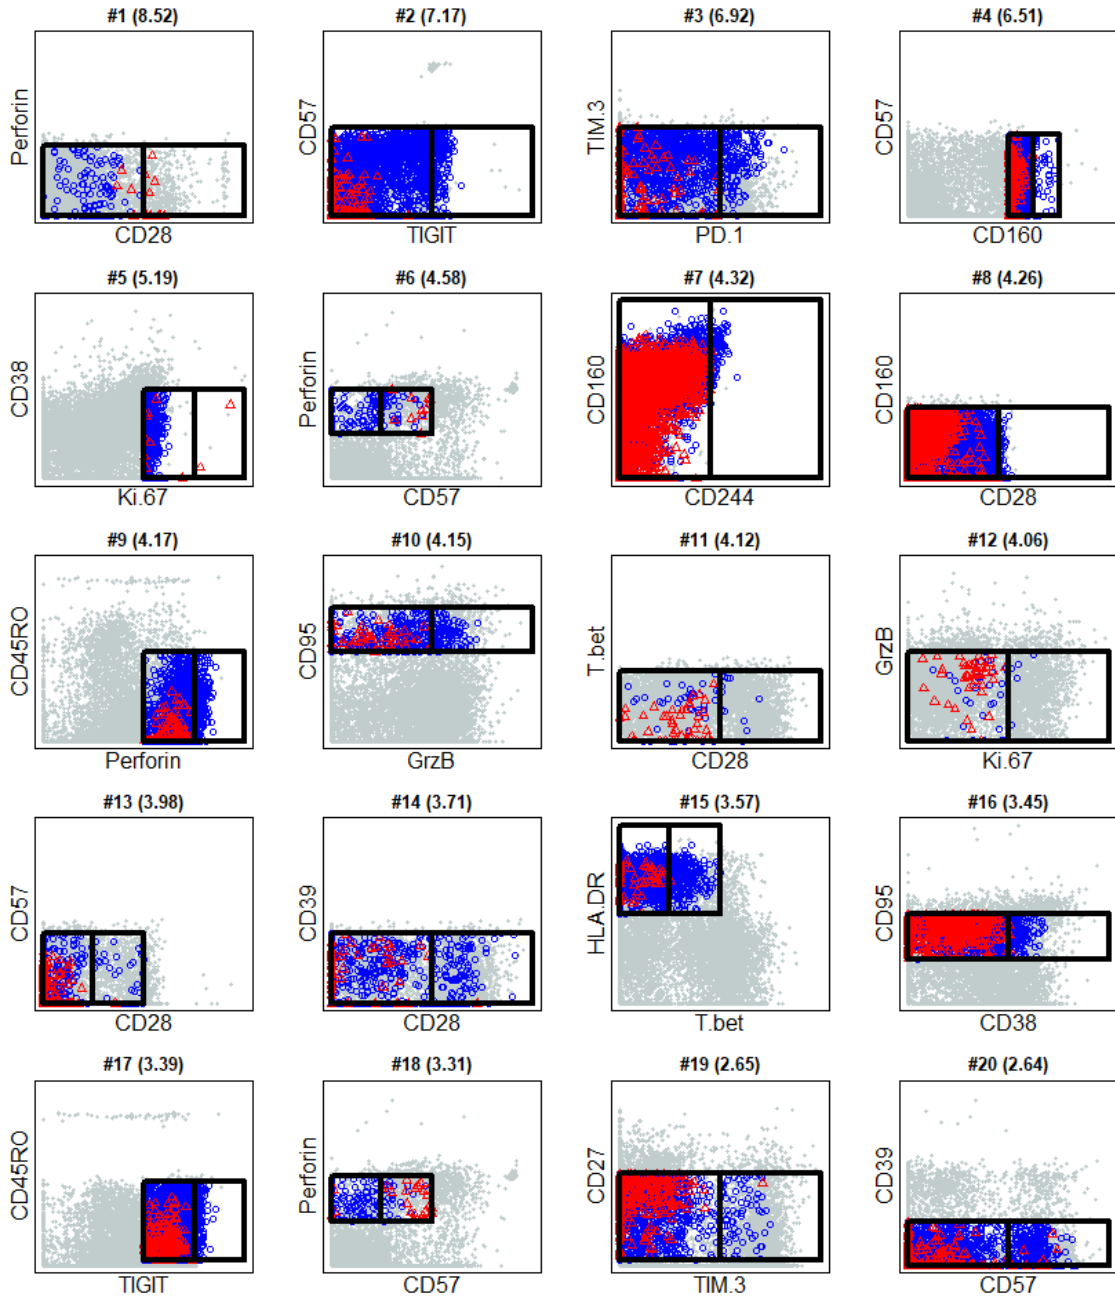


Figure 14: The solid lines delineate the nodes with the highest values of $\text{eff}(A)$ and their two children. The red triangle points and the blue circle points are the observations from the two samples in the node. The observations outside the node are in gray. In this figure, the nodes with $n_g(A) \geq 10$ ($g = 1, 2$) are chosen.