Exact Non-Parametric Bayesian Inference on Infinite Trees

Marcus Hutter

RSISE @ ANU and SML @ NICTA
Canberra, ACT, 0200, Australia*
marcus@hutter1.net www.hutter1.net

31 March 2009

Abstract

Given i.i.d. data from an unknown distribution, we consider the problem of predicting future items. An adaptive way to estimate the probability density is to recursively subdivide the domain to an appropriate data-dependent granularity. In Bayesian inference one assigns a data-independent prior probability to "subdivide", which leads to a prior over infinite(ly many) trees. We derive an exact, fast, and simple inference algorithm for such a prior, for the data evidence, the predictive distribution, the effective model dimension, moments, and other quantities. We prove asymptotic convergence and consistency results, and illustrate the behavior of our model on some prototypical functions.

Contents

Keywords

Bayesian density estimation, exact linear time algorithm, non-parametric inference, adaptive infinite tree, Polya tree, scale invariance, consistency, asymptotics.

*Preliminary results have been presented at the AISTATS 2005 conference [Hut05a].

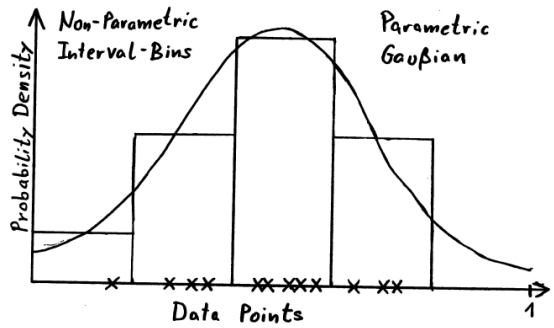

Figure 1: Bins versus Gaussian estimate of the true=data-generating probability density. (More decent diagrams will be made for the final version).

1 Introduction

Inference. We consider the problem of inference from i.i.d. data D, in particular of the unknown distribution q the data is sampled from. In case of a continuous domain this means inferring a probability density from data. Without structural assumption on q, this is hard to impossible, since a finite amount of data is never sufficient to uniquely select a density (model) from an infinite-dimensional space of densities (model class).

Methods. In parametric estimation one assumes that q belongs to a finite-dimensional family. The two-dimensional family of Gaussians characterized by mean and variance is prototypical (Figure 1). The maximum likelihood (ML) estimate of q is the distribution that maximizes the data likelihood. Maximum likelihood overfits if the family is too large and especially if it is infinite-dimensional. A remedy is to penalize complex distributions by assigning a prior (2nd order) probability to the densities q. Maximizing the model posterior (MAP), which is proportional to likelihood times the prior, prevents overfitting. A full Bayesian procedure keeps the complete posterior for inference. Typically, summaries like the mean and variance of the posterior are reported.

How to choose the prior? In finite or small compact low-dimensional spaces a uniform prior often works (MAP reduces to ML). In the non-parametric case one typically devises a hierarchy of finite-dimensional model classes of increasing dimension. Selecting the dimension with maximal posterior often works well due to the Bayes factor phenomenon [Goo83, Goo84, Jef35, Jay03, Mac03]: In case the true model is low-dimensional, higher-dimensional (complex) model classes are automatically penalized, since they contain fewer "good" models. In a full Bayesian treatment one would assign a prior probability (e.g. $\frac{1}{d^2}$) to the dimension d and mix over the dimension.

Interval Bins. The probably simplest and oldest model for an interval domain is to divide the interval (uniformly) into bins, assume a constant distribution within

each bin, and take a frequency estimate for the probability in each bin (Figure 1), or a Dirichlet posterior in Bayesian inference. There are heuristics for choosing the number of bins as a function of the data size. The simplicity and easy computability of the bin model is very appealing to practitioners. Drawbacks are that distributions are discontinuous, its restriction to one dimension (or at most low dimension: curse of dimensionality), the uniform (or more generally fixed) discretization, and the heuristic choice of the number of bins. We present a full Bayesian solution to these problems, except for the non-continuity problem. Our model can be regarded as an extension of Polya trees [Fer73, Lav92, Lav94].

Related work. There are plenty of alternative Bayesian models that overcome some or all of the limitations. Examples are continuous Dirichlet process (mixtures) [Fer73], Bernstein polynomials [PW02], Bayesian field theory [Lem03], randomized Polya trees [PRLW03], Bayesian bins with boundary averaging [EF05], Bayesian kernel density estimation or other mixture models [EW95], and universal priors [Hut05b], but exact analytical solutions are infeasible. Markov Chain Monte Carlo sampling [Bis06], Expectation Maximization algorithms [DLR77], variational methods [Bis06], efficient MAP or M(D)L approximations [KM07], or kernel density estimation [GM03] can often be used to obtain approximate numerical solutions, but computation time and/or global convergence remain critical issues. There are of course also plenty of non-Bayesian density estimators; see (references in) [KF98, BM98, LLW07] in general, and [KK97, KF98] for density tree estimation in particular.

Our tree mixture model. The idea of the model class discussed in this paper is very simple: With some (e.g. equal) probability, we chose q either uniform or split the domain in two parts (of equal volume), and assign a prior to each part, recursively, i.e. in each part again either uniform or split. For finitely many splits, q is a piecewise constant function, for infinitely many splits it is virtually *any* distribution. While the prior over q is neutral about uniform versus split, we will see that the posterior favors a split if and only if the data clearly indicates non-uniformity. The method is a full Bayesian non-heuristic tree approach to adaptive binning for which we present a very simple and fast algorithm for computing all(?) quantities of interest.

Note that we are not arguing that our model performs better in practice than the more advanced models above. The main distinguishing feature of our model is that it allows for a fast and exact analytical solution. It's likely use is as a building block in complex problems, where computation time and Bayesian integration are the major issues. In any case, if/since the Polya tree model deserves attention, also our model should.

Contents. In Section 2 we introduce our model and compare it to Polya trees. We also discuss some example domains, like intervals, strings, volumes, and classification tasks. Section 3 derives recursions for the posterior and the data evidence. Section 4 proves convergence/consistency. In Section 5 we introduce further quantities of interest, including the effective model dimension, the tree size and height, the cell

volume, and moments, and present recursions for them. The proper case of infinite trees is discussed in Section 6, where we analytically solve the infinite recursion at the data separation level. Section 7 collects everything together and presents the algorithm. In Section 8 we numerically illustrate the behavior of our model on some prototypical functions. Section 9 contains a brief summary, conclusions, and outlook, including natural generalizations of our model. See [Hut07] for program code.

2 The Tree Mixture Model

Setup and basic quantities of interest. We are given i.i.d. data $D=(x^1,...,x^n)\in\Gamma^n$ of size n from domain Γ, e.g. $\Gamma\subseteq\mathbb{R}^d$, sampled from some unknown probability density $q:\Gamma\to\mathbb{R}$. Standard inference problems are to estimate q from D or to predict the next data item $x^{n+1}\in\Gamma$. By definition, the (objective or aleatoric) data likelihood density under model q is

$$\text{likelihood:}\qquad p(D|q) \;\equiv\; q(x_1)\cdot...\cdot q(x_n) \tag{1}$$

Note that we consider sorted data, which avoids annoying multinomial coefficients. Otherwise this has no consequences. Results are independent of the order and depend on the counts only, as they should. A Bayesian assumes a (belief or 2^{nd}-order or epistemic or subjective) prior over models q in some model class Q:

$$\text{prior:}\qquad p(q)\quad\text{with}\quad q\in Q$$

The data evidence is

$$\text{evidence:}\qquad p(D) \;=\; \int_Q p(D|q)p(q)dq \tag{2}$$

Having the evidence, Bayes' famous rule allows to compute the (belief or 2^{nd}-order or epistemic or subjective) posterior of q:

$$\text{posterior:}\qquad p(q|D) \;=\; \frac{p(D|q)p(q)}{p(D)} \tag{3}$$

The predictive distribution, i.e. the conditional probability that next data item is $x=x^{n+1}$, given D, follows from the evidences of D and (D,x):

$$\text{predictive distribution:}\qquad p(x|D) \;=\; \frac{p(D,x)}{p(D)} \tag{4}$$

Since the posterior is a complex object, we need summaries like the expected q-probability of x and (co)variances. Fortunately they can also be reduced to computation of evidences:

$$\begin{aligned} E[q(x)|D] \;&:=\; \int q(x)p(q|D)dq \;=\; \int q(x)\frac{p(D|q)p(q)}{p(D)}dq \\ &=\; \frac{\int p(D,x|q)p(q)dq}{p(D)} \;=\; \frac{p(D,x)}{p(D)} \;=\; p(x|D) \end{aligned}$$

where we used the formulas for the posterior, the likelihood, the evidence, and the predictive distribution, in this order. Similarly for the covariance we obtain

$$\begin{aligned}&\mathrm{Cov}[q(x)q(y)|D]\\ \equiv\ & E[q(x)q(y)|D] - E[q(x)|D]\cdot E[q(x)|D]\\ =\ & p(x,y|D) - p(x|D)p(y|D)\end{aligned}$$

We derive and discuss further summaries of q for our particular tree model, like the model complexity or effective dimension, the tree height or cell size, and moments later.

Hierarchical tree partitioning. So far everything has been fairly general. We now introduce the tree representation of domain Γ. We partition Γ into Γ_0 and Γ_1, i.e. $\Gamma=\Gamma_0\cup\Gamma_1$ and $\Gamma_0\cap\Gamma_1=\emptyset$. Recursively we (sub)partition $\Gamma_z=\Gamma_{z0}\dot\cup\Gamma_{z1}$ for $z\in\mathbb{B}_0^m$, where $\mathbb{B}_k^m:=\bigcup_{i=k}^m\{0,1\}^i$ is the set of all binary strings of length between k and m, and $\Gamma_\epsilon=\Gamma$, where $\epsilon=\{0,1\}^0$ is the empty string. We are interested in an infinite recursion, but for convenience we assume a finite tree height $m<\infty$ and consider $m\to\infty$ later. Also let $l:=\ell(z)$ be the length of string $z=z_1...z_l=:z_{1:l}$, and $|\Gamma_z|$ the volume or length or cardinality of Γ_z.

Example spaces (Figures 2 & 3). *Intervals:* Assume $\Gamma=[0,1)$ is the unit interval, recursively bisected into intervals $\Gamma_z=[0.z,0.z+2^{-l})$ of length $|\Gamma_z|=2^{-l}$, where $0.z$ is the real number in $[0,1)$ with binary expansion $z_1...z_l$.

Strings: Assume $\Gamma_z=\{zy:y\in\{0,1\}^{m-l}\}$ is the set of strings of length m starting with z. Then $\Gamma=\{0,1\}^m$ and $|\Gamma_z|=2^{m-l}$. For $m=\infty$ this set is continuous, for $m<\infty$ finite.

Trees: Let Γ be a complete binary tree of height m and Γ_{z0} (Γ_{z1}) be the left (right) subtree of Γ_z. If $|\Gamma_z|$ is defined as one more than the number of nodes in Γ_z, then $|\Gamma_z|=2^{m+1-l}$.

Volumes: Consider $\Gamma\subset\mathbb{R}^d$, e.g. the hypercube $\Gamma=[0,1)^d$. We recursively halve Γ_z with a hyperplane orthogonal to dimension (l mod d)+1, i.e. we sweep through all orthogonal directions. $|\Gamma_z|=2^{-l}|\Gamma|$.

Compactification: We can compactify $\Gamma\subseteq(1,\infty]$ (this includes $\Gamma=\mathbb{N}\setminus\{1\}$) to the unit interval $\Gamma':=\{\frac{1}{x}:x\in\Gamma\}\subseteq[0,1)$, and similarly $\Gamma\subseteq\mathbb{R}$ (this includes $\Gamma=\mathbb{Z}$) to $\Gamma':=\{x\in[0,1):\frac{2x-1}{x(1-x)}\in\Gamma\}$. All reasonable spaces can be reduced to one of the spaces described above, although this reduction may introduce unwanted artifacts.

Classification: Consider an observation $o\in\Gamma'$ (e.g. email) that is classified as $c\in\{0,1\}$ (e.g. good versus spam), where Γ' could be one of the spaces above (e.g. o is a sequence of binary features in decreasing order of importance). Then $x:=(o,c)\in\Gamma:=\Gamma'\times\{0,1\}$ and $\Gamma_{0z}=\Gamma'_z\times\{0\}$ and $\Gamma_{1z}=\Gamma'_z\times\{1\}$. Given D (e.g. pre-classified emails), a new observation o is classified as c with probability $p(c|D,o)\propto p(D,x)$. Similar for more than two classes.

In all these examples we have (chosen) $|\Gamma_{z0}|=|\Gamma_{z1}|=\frac{1}{2}|\Gamma_z|\ \forall z\in\mathbb{B}_0^{m-1}$, and this is the only property we need and henceforth assume. W.l.g. we assume/define/

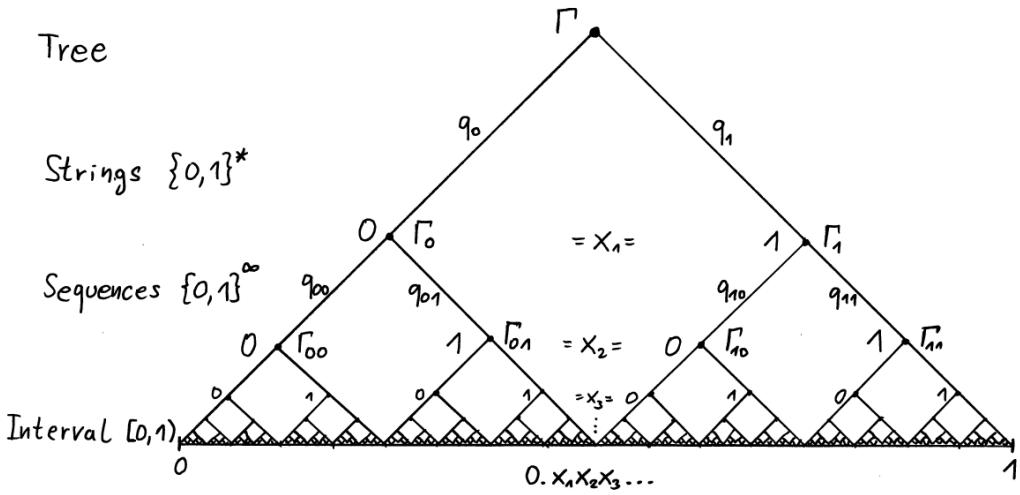

Figure 2: Tree representation of strings or sequences or intervals.

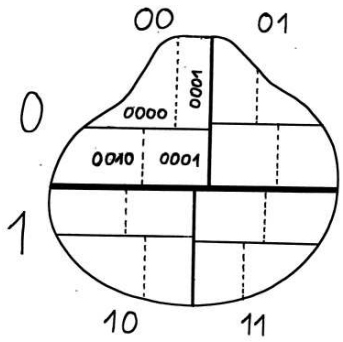

Figure 3: Tree representation of volumes.

rescale $|\Gamma| = 1$. Generalizations to non-binary and non-symmetric partitions are straightforward and briefly discussed at the end.

Identification. We assume that $\{\Gamma_z : z \in \mathbb{B}_0^m\}$ are (basis) events that generate our σ-algebra. For every $x \in \Gamma$ let x' be the string of length $\ell(x') = m$ such that $x \in \Gamma_{x'}$. We assume that distributions q are σ-measurable, i.e. to be constant on $\Gamma_{x'}$ $\forall x' \in \mathbb{B}^m$. For $m = \infty$ this assumption is vacuous; we get *all* Borel measures. Hence, we can identify the continuous sample space Γ with the (for $m < \infty$ discrete) space $\mathbb{B}^m$ of binary sequences of length m, i.e. in a sense all example spaces are isomorphic. While we have the volume model in mind for real-world applications, the string model will be convenient for mathematical notation, the tree metaphor will be convenient in discussion, and the interval model will be easiest to implement and to present graphically.

Notation. As described above, Γ may also be a tree. This interpretation suggests the following scheme for defining the probability of q on the leaves x'. The probability of the left child node $z0$, given we are in the parent node z, is $P[\Gamma_{z0}|\Gamma_z,q]$, so we have

$$p(x|\Gamma_z, q) = p(x|\Gamma_{z0}, q) \cdot P[\Gamma_{z0}|\Gamma_z, q] \quad \text{if} \quad x \in \Gamma_{z0}$$

and similarly for the right child. In the following we often have to consider distributions conditioned to and in the subtree Γ_z, so the following notation will turn out convenient

$$q_{z0} := P[\Gamma_{z0}|\Gamma_z, q], \quad q_{z1} := P[\Gamma_{z1}|\Gamma_z, q], \quad p_z(x|...) := 2^{-l} p(x|\Gamma_z ...) \tag{5}$$
$$\Rightarrow p_z(x|q) = 2q_{zx_{l+1}} p_{zx_{l+1}}(x|q) = ... = \prod_{i=l+1}^{m} 2q_{x_{1:i}} \text{ if } x \in \Gamma_z$$

where $p(x|\Gamma_{x'},q) := |\Gamma_{x'}|^{-1} = 2^m$ is uniform (by assumption). Note that $q_{z0} + q_{z1} = 1$. Finally, let

$$\vec{q}_{z*} \;:=\; (q_{zy} : y \in \mathbb{B}_1^{m-l})$$

be the $(2^{m-l+1}-2)$-dimensional *vector* or *ordered set* or *tree* of all *reals* $q_{zy} \in [0,1]$ in subtree Γ_z. Note that $q_z \notin \vec{q}_{z*}$. The *(non)density* $q_z(x) := p_z(x|q)$ depends on all and only these q_{zy}. For $z \neq \epsilon$, $q_z()$ and $p_z()$ are only proportional to a density due to the factor 2^{-l}, which has been introduced to make $p_{x'}(x|...) \equiv 1$. (They are densities w.r.t. $2^l \lambda_{|\Gamma_z}$, where λ is the Lebesgue measure.) We have to keep this in mind in our derivations, but can ignore this widely in our discussion.

Polya trees. In the Polya tree model one assumes that the $q_{z0} \equiv 1 - q_{z1}$ are independent and Beta$(\cdot,\cdot)$ distributed, which defines the prior over q. Polya trees form a conjugate prior class, since the posterior is also a Polya tree, with empirical counts added to the Beta parameters. If the same Beta is chosen in each node, the posterior of x is pathological for $m \to \infty$: The density does nowhere exist with probability 1. A cure is to increase the Beta parameters with l, e.g. quadratically, but this results in "underfitting" for large sample sizes, since Beta(large,large) is too informative and strongly favors q_{z0} near $\frac{1}{2}$. It also violates scale invariance, which should ideally

hold if we do not have any prior knowledge about the scale. That is, the p(oste)rior in $\Gamma_0=[0,\frac{1}{2})$ should be the same as for $\Gamma=[0,1)$ (after rescaling all $x\rightsquigarrow x/2$ in D).

The new tree mixture model. The prior $p(q)$ follows from specifying a prior over $\vec{q}_*$, since $q(x)\propto q_{x_1}\cdot...\cdot q_{x_{1:m}}$ by (5). The distribution in each subset $\Gamma_z\subseteq\Gamma$ shall be either *u*niform or non-uniform. A necessary (but not sufficient) condition for uniformity is $q_{z0}=q_{z1}=\frac{1}{2}$.

$$p^u(q_{z0},q_{z1}) \;:=\; \delta(q_{z0}-\tfrac{1}{2})\delta(q_{z1}-\tfrac{1}{2}), \tag{6}$$

where $\delta()$ is the Dirac delta. To get uniformity on Γ_z we have to recurse the tree down in this way.

$$p^u_z(\vec{q}_{z*}) \;:=\; p^u(q_{z0},q_{z1})p^u_{z0}(\vec{q}_{z0*})p^u_{z1}(\vec{q}_{z1*}) \tag{7}$$

with the natural recursion termination $p^u_z(\vec{q}_{z*})=1$ when $\ell(z)=m$, since then $\vec{q}_{z*}=\emptyset$. For a non-uniform distribution on Γ_z we allow any probability split $q(\Gamma_z)=q(\Gamma_{z0})+q(\Gamma_{z1})$, or equivalently $1=q_{z0}+q_{z1}$. We assume a Beta prior on the *s*plit. Scale invariance requires the Beta parameters to be the same in all nodes of the tree and symmetry demands a symmetric Beta, i.e.

$$\begin{aligned} p^s(q_{z0},q_{z1}) &:= \mathrm{Beta}(q_{z0},q_{z1}|\alpha,\alpha), &(8)\\ \mathrm{Beta}(p,q|\alpha,\beta) &:= \tfrac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}p^{\alpha-1}q^{\beta-1}\delta(p+q-1) &(9)\end{aligned}$$

where $\Gamma(\alpha)=\int_0^\infty t^{\alpha-1}\mathrm{e}^{-t}dt$ is the Gamma function. For $\alpha=1$ this specializes to the natural uniform prior $p^s(q_{z0},q_{z1})=\delta(q_{z0}+q_{z1}-1)$ on the *s*plit. We now recurse down the tree

$$p^s_z(\vec{q}_{z*}) \;:=\; p^s(q_{z0},q_{z1})p_{z0}(\vec{q}_{z0*})p_{z1}(\vec{q}_{z1*}) \tag{10}$$

again with the natural recursion termination $p_z(\vec{q}_{z*})=p(\emptyset)=1$ when $\ell(z)=m$. Finally we have to mix the uniform with the non-uniform case.

$$p_z(\vec{q}_{z*}) \;:=\; u\cdot p^u_z(\vec{q}_{z*}) + s\cdot p^s_z(\vec{q}_{z*}) \tag{11}$$

with $u,s\in[0,1]$ and $u+s=1$. The 50/50 mixture $u=s=\frac{1}{2}$ will be of special interest. This completes the specification of the prior $p(q)=p_\epsilon(\vec{q}_*)$.[1]

For example, if the first bit in x is a class label and the remaining are binary features in decreasing order of importance, then given class and features $z=x_{1:l}$, further features $x_{l+1:m}$ could be relevant for classification ($q_z(x)$ is non-uniform) or irrelevant ($q_z(x)$ is uniform).

Comparison to the Polya tree. Note the important difference in the recursions (7) and (10). Once we decided on a uniform distribution (6) we have to equally split probabilities down the recursion to the end, i.e. we recurse in (7) with p^u, rather than the mixture p (this actually allows to solve the recursion). On the other hand if

[1]Note that $p_z(\vec{q}_{z*})$ is *not* the marginal of $p(q)$ to $\vec{q}_{z*}$, but one can show that $p_z(\vec{q}_{z*})=p(\vec{q}_{z*}|q_{z_1}\neq\frac{1}{2},...,q_{z_{1:l}}\neq\frac{1}{2})$ and optionally additional conditions on some or all $q\notin\vec{q}_{z*}$.

we decided on a non-uniform split (8), the left and right partition each itself may be uniform or not, i.e. we recurse in (10) with the mixture p, rather than p^s. Inserting (8) in (10) in (11) and recursively (6) in (7) in (11) we get the following recursion for the prior

$$p_z(\vec{q}_{z*}) = u\cdot\prod_{y\in\mathbb{B}_1^{m-l}}\delta(q_{zy}-\tfrac{1}{2}) + s\cdot\text{Beta}(q_{z0},q_{z1}|\alpha,\alpha)p_{z0}(\vec{q}_{z0*})p_{z1}(\vec{q}_{z1*}) \tag{12}$$

Choosing $u=0$ would lead to the Polya tree model (and its problems) with $q_{z0}\sim$ Beta$(\cdot|\alpha,\alpha)$. With p instead of p^u on the r.h.s. of (7) we would get a quasi-Polya model (same problems) with $q_{z0}\sim u\cdot\text{Beta}(\cdot|\infty,\infty)+s\cdot\text{Beta}(\cdot|\alpha,\alpha)$.

For $m\to\infty$, our model is "scale" invariant *and* leads to continuous densities for $n\to\infty$, unlike the Polya tree model. We also don't have to tune Beta parameters. We can use a non-informative prior like $\alpha=1$ and $u=s=\frac{1}{2}$. The model "tunes itself" by suitably assigning high/low posterior probability to subdividing cells. While Polya trees form a natural conjugate prior class, our prior does not directly, but may be generalized to do so. The computational complexity for the quantities of interest will be the same (essentially $O(n)$), i.e. as good as it could be.

Formal and effective dimension. Formally our model is $2\cdot(2^m-1)$-dimensional, but the effective dimension can by much smaller, since $\vec{q}_*$ is forced with a non-zero probability to a much smaller polytope, for instance with probability u to the zero-dimensional globally uniform distribution. We will compute the effective p(oste)rior dimension. Alternatively, we could have considered a mixture over all ($\hat{=}$ lower dimensional) partial trees with Γ_z as leaf if q is uniform on Γ_z, but considering one complete tree is more convenient for analytical manipulation.

3 Evidence and Posterior Recursion

At the end of Section 2 we defined our tree mixture model. The next step is to compute the standard quantities of interest defined at the beginning of Section 2. The evidence (2) is key, the other quantities (posterior, predictive distribution, expected $q(x)$ and its variance) follow then immediately. Let

$$D_z \;:=\; \{x\in D : x\in\Gamma_z\}$$

be the $n_z:=|D_z|$ data points that lie in subtree Γ_z. We compute $p_z(D_z)$ recursively for all $z\in\mathbb{B}_0^{m-1}$, which gives $p(D)=p_\epsilon(D_\epsilon)$.

Theorem 1 (Evidence recursion) *For $\ell(z)<m$ the recursion for the evidence is*

$$p_z(D_z) \;=\; u \;+\; s\cdot\frac{p_{z0}(D_{z0})p_{z1}(D_{z1})}{w(n_{z0},n_{z1})} \tag{13}$$

$$w(n_{z0},n_{z1}) := \frac{2^{-n_z}\cdot\Gamma(n_z+2\alpha)}{\Gamma(n_{z0}+\alpha)\Gamma(n_{z1}+\alpha)}\cdot\frac{\Gamma(\alpha)^2}{\Gamma(2\alpha)} =: w_{n_z}(\Delta_z) \tag{14}$$

$$n_z \;=\; n_{z0}+n_{z1}, \quad \Delta_z \;:=\; \frac{n_{z0}}{n_z}-\tfrac{1}{2}$$

The recursion terminates with $p_z(D_z)\equiv 1$ when $\ell(z)=m$.

The recursion (13) follows by multiplying (16) in Theorem 2 (stated and proven below) with $p_z(D_z)$ and adding u. For $\alpha=1$, (13) and in particular the weight $w_{n_z}=2^{-n_z}(n_z+1)\binom{n_z}{n_{z0}}$ can be interpreted as follows: With probability u, the evidence is uniform in Γ_z. Otherwise data D_z is split into two partitions of size n_{z0} and $n_{z1}=n_z-n_{z0}$. First, choose n_{z0} uniformly in $\{0,...,n_z\}$. Second, given n_z, choose uniformly among the $\binom{n_z}{n_{z0}}$ possibilities of selecting n_{z0} out of n_z data points for Γ_{z0} (the remaining n_{z1} are then in Γ_{z1}). Third, distribute D_{z0} according to $p_{z0}(D_{z0})$ and D_{z1} according to $p_{z1}(D_{z1})$. Then, the evidence in case of a split is the second term in (13). The factor 2^{n_z} is due to our normalization convention (5). This also verifies that the r.h.s. yields the l.h.s. if integrated over all D_z, as it should be. For $n_z\to\infty$ we will show in Section 4 that $w_{n_z}\to\infty$ if $n_{z0}\sim n_{z1}\to\infty$ and $w_{n_z}\to 0$ otherwise, indicating that the weight w is large (small) for (non)uniform distribution, as it should be.

Theorem 2 (Posterior recursion) *For $\ell(z)<m$ the recursion for the posterior is*

$$\begin{aligned} p_z(\vec{q}_{z*}|D_z) &= \frac{u}{p_z(D_z)}\prod_{y\in\mathbb{B}_1^{m-l}}\delta(q_{zy}-\tfrac{1}{2}) \qquad (15)\\ &+ g_z(D_z)\mathrm{Beta}(q_{z0},q_{z1}|n_{z0}+\alpha,n_{z1}+\alpha)p_{z0}(\vec{q}_{z0*}|D_{z0})p_{z1}(\vec{q}_{z1*}|D_{z1})\\ g_z(D_z) &:= s\cdot\frac{p_{z0}(D_{z0})p_{z1}(D_{z1})}{p_z(D_z)w(n_{z0},n_{z1})} = 1-\frac{u}{p_z(D_z)} \qquad (16)\end{aligned}$$

The recursion terminates with $p_z(\vec{q}_{z*}|D_z)\equiv 1$ when $\ell(z)=m$.

$g_z(D_z)$ may be interpreted as the posterior probability of splitting Γ_z.

Proof. Using Bayes rule (3) we represent the posterior as

$$p_z(\vec{q}_{z*}|D_z)p_z(D_z) = p_z(D_z|\vec{q}_{z*})p_z(\vec{q}_{z*}) \qquad (17)$$

and further substitute $p_z(\vec{q}_{z*})=up_z^u(\vec{q}_{z*})+sp_z^s(\vec{q}_{z*})$ (11). For the uniform part we get

$$\begin{aligned} p_z(D_z|\vec{q}_{z*})\cdot p_z^u(\vec{q}_{z*}) &= \prod_{x\in D_z}(2q_{x_{1:l+1}}\cdot....\cdot 2q_{x_{1:m}})\cdot\prod_{y\in\mathbb{B}_1^{m-l}}\delta(q_{zy}-\tfrac{1}{2})\\ &= \prod_{y\in\mathbb{B}_1^{m-l}}\delta(q_{zy}-\tfrac{1}{2}), \qquad (18)\end{aligned}$$

where we recursively inserted (6) in (7), and (1) and (5) into (18). Due to the δ, we can simply set all $q_{zy}=\frac{1}{2}$. For the split we get

$$\begin{aligned} &p_z(D_z|\vec{q}_{z*})\cdot p_z^s(\vec{q}_{z*})\\ &= \Big(\prod_{x\in D_{z0}}2q_{z0}\Big)p_{z0}(D_{z0}|\vec{q}_{z0*})\Big(\prod_{x\in D_{z1}}2q_{z1}\Big)p_{z1}(D_{z1}|\vec{q}_{z1*}) \qquad (19)\end{aligned}$$

$$
\begin{aligned}
& \times \tfrac{\Gamma(2\alpha)}{\Gamma(\alpha)^2}\, q_{z0}^{\alpha-1} q_{z1}^{\alpha-1} \delta(q_{z0}+q_{z1}-1) p_{z0}(\vec{q}_{z0*}) p_{z1}(\vec{q}_{z1*}) \\
= \;& 2^{n_z} \tfrac{\Gamma(2\alpha)}{\Gamma(\alpha)^2}\, q_{z0}^{n_{z0}+\alpha-1} q_{z1}^{n_{z1}+\alpha-1} \delta(q_{z0}+q_{z1}-1) \qquad (20) \\
& \times p_{z0}(\vec{q}_{z0*}|D_{z0}) p(D_{z0}) p_{z1}(\vec{q}_{z1*}|D_{z1}) p(D_{z1}) \\
= \;& \tfrac{1}{s} g_z(D_z) p_z(D_z) \mathrm{Beta}(q_{z0}, q_{z1}|n_{z0}+\alpha, n_{z1}+\alpha) \qquad (21) \\
& \times p_{z0}(\vec{q}_{z0*}|D_{z0}) p_{z1}(\vec{q}_{z1*}|D_{z1})
\end{aligned}
$$

In (19) we split D_z into D_{z0} and D_{z1} and used (1) and (5) and the fact that $p_{z0}(D_{z0}|\vec{q}_{z*})$ depends on q through $\vec{q}_{z0*}$ only. We also inserted (9) in (8) in (10) in (19) and used $n_{z0}+n_{z1}=n_z$. Rearranging terms and using Bayes rule (17) for subtrees Γ_{z0} and Γ_{z1} we get (20). The last equality is easiest proven backwards by inserting g_z (16) and w (14) and Beta (9) into (21). Inserting (11) and (18) and (19)-(21) into (17) and dividing by $p_z(D_z)$ yields (15).

Integrating (15) over $\vec{q}_{z*}$ and noting that $\int d\vec{q}_{z*} = \int dq_{z0} dq_{z1} \cdot \int d\vec{q}_{z0*} \cdot \int d\vec{q}_{z1*}$ factorizes and that $\prod\delta()$ and Beta() and the $\vec{q}_{z*}$, $\vec{q}_{z0*}$ and $\vec{q}_{z1*}$ posteriors are all proper densities which integrate to 1, we get

$$
1 \;=\; \frac{u}{p_z(D_z)} \cdot 1 + g_z(D_z) \cdot 1 \cdot 1 \cdot 1
$$

This shows the last equality in (16). Theorem 1 (13) now follows by multiplying (16) with $p_z(D_z)$ and adding u.

For a formal proof of the recursion termination, recall (5): For $\ell(z)=m$ and $x \in \Gamma_z$ we have $\Gamma_{x'}=\Gamma_z \Rightarrow p_z(x|q)=1 \Rightarrow p_z(D_z|q)=1 \Rightarrow p_z(D_z)=1$. □

4 Asymptotic Convergence/Consistency ($n \to \infty$)

Discussing the weight. The relative probability of splitting (second term on r.h.s. of (13)) to the uniform case (first term in r.h.s. of (13)) is controlled by the weight w. Large (small) weight indicates a (non)uniform distribution, provided p_{z0} and p_{z1} are $O(1)$. The balance $\Delta_z \approx 0$ ($\not\approx 0$) indicates a (non)symmetric partitioning of the data among the left and right branch of Γ_z. Asymptotically for large n_z (and small Δ_z), we have

$$
w_{n_z}(\Delta_z) \approx c_\alpha \sqrt{\tfrac{2n_z}{\pi}}\, \mathrm{e}^{-2n_z\Delta_z^2}
$$

where $c_\alpha > 0$ is some finite constant. Assume that data D is sampled from the true distribution $\dot{q}$. The probability of the left branch Γ_{z0} of Γ_z is $\dot{q}_{z0} \equiv P[\Gamma_{z0}|\Gamma_z,\dot{q}] = 2^l \dot{q}_z(\Gamma_{z0})$. The relative frequencies $\frac{n_{z0}}{n_z}$ asymptotically converge to $\dot{q}_{z0}$. More precisely $\frac{n_{z0}}{n_z} = \dot{q}_{z0} \pm O(n_z^{-1/2})$. Similarly for the right branch. Assume the probabilities are equal ($\dot{q}_{z0}=\dot{q}_{z1}=\frac{1}{2}$), possibly but not necessarily due to a uniform $\dot{q}_z()$ on Γ_z. Then $\Delta_z = O(n_z^{-1/2})$, which implies

$$
w_{n_z}(\Delta_z) \approx \Theta(\sqrt{n}_z) \;\stackrel{n_z\to\infty}{\longrightarrow}\; \infty \quad \text{if} \quad \dot{q}_{z0} = \dot{q}_{z1} = \tfrac{1}{2},
$$

consistent with our anticipation. Conversely, for $\dot{q}_{z0} \neq \dot{q}_{z1}$ (which implies non-uniformity of $\dot{q}_z()$) we have $\Delta_z \to c := \dot{q}_{z0} - \frac{1}{2} \neq 0$, which implies

$$w_{n_z}(\Delta_z) \approx \sqrt{\tfrac{2n_z}{\pi}}\, \mathrm{e}^{-2n_z c^2} \stackrel{n_z\to\infty}{\longrightarrow} 0 \quad \text{if} \quad \dot{q}_{z0} \neq \dot{q}_{z1},$$

again, consistent with our anticipation. Formally, the following can be proven:

Theorem 3 (Weight asymptotics) *For $\dot{q}_{z0} = \frac{1}{2}$ we have with probability 1 (w.p.1)*

$$i) \quad \lim_{n_z\to\infty} \frac{\ln n_z}{\sqrt{n_z}}\, w_{n_z}(\Delta_z) = \infty, \quad \textit{and}$$

$$ii) \quad \limsup_{n_z\to\infty} \sqrt{\frac{\pi}{2n_z}}\, w_{n_z}(\Delta_z) = c_\alpha \;\; \{ \begin{smallmatrix} >0 \\ <\infty \end{smallmatrix}$$

where $c_\alpha = 4^{\alpha-1}\Gamma(\alpha)^2/\Gamma(2\alpha)$. For $\dot{q}_{z0} \neq \frac{1}{2}$ we have w.p.1

$$iii) \quad \lim_{n_z\to\infty} \mathrm{e}^{2n_z c^2} w_{n_z}(\Delta_z) = 0 \quad \forall\, |c| < |\dot{q}_{z0} - \tfrac{1}{2}|$$

Proof. We will drop the index z everywhere. We need an asymptotic representation of w for $n_0, n_1 \to \infty$. Using Stirling's approximation $\ln\Gamma(x) = (x-\frac{1}{2})\ln x - x + \frac{1}{2}\ln(2\pi) + O(\frac{1}{x})$ we get after some algebra

$$\begin{aligned}
\ln w_n(\Delta) &= -n[H(\tfrac{1}{2}) - H(\tfrac{1}{2} \pm \tilde{\Delta})] + \tfrac{1}{2}\ln\tfrac{n}{2\pi} \\
&\quad + (2\alpha-1)H(\tfrac{1}{2} \pm \tilde{\Delta}) + C_\alpha + O(\tfrac{\alpha}{n_0} + \tfrac{\alpha}{n_1}), \\
H(p) &= -p\ln p - (1-p)\ln(1-p) \;=\; \text{Entropy}(p), \\
\tilde{\Delta} &= \tfrac{n}{n+2\alpha-1}\Delta_n \;=\; \tfrac{\frac{1}{2}(n_0-n_1)}{n+2\alpha-1}, \quad \Delta_n = \Delta = \tfrac{n_0}{n} - \tfrac{1}{2}, \\
C_\alpha &= 2\ln\Gamma(\alpha) - \ln\Gamma(2\alpha)
\end{aligned} \tag{22}$$

(i) follows from the law of the iterated logarithm

$$\limsup_{n\to\infty} \frac{|X_1 + ... + X_n - n\mu|}{\sigma\sqrt{n \ln\ln n}} \;=\; 1 \quad \text{w.p.1}$$

for i.i.d. random variables $X_1,...,X_n$ with mean μ and variance σ^2. For the i^{th} data item in D, let $X_i = 1$ if $x \in D_0$ and $X_i = 0$ if $x \in D_1$. Then the X_i are i.i.d. Bernoulli($\dot{q}_0$) with $\mu = \dot{q}_0 = \frac{1}{2}$ and $\sigma^2 = \dot{q}_0\dot{q}_1 = \frac{1}{4}$. Further, $X_1 + ... + X_n = n_0$ implies $X_1 + ... + X_n - n\mu = n\Delta_n$ implies $\limsup_n \sqrt{\frac{4n}{\ln\ln n}}|\Delta_n| = 1$ w.p.1. implies

$$\Delta_n^2 \;\leq\; (1+\varepsilon)\tfrac{\ln\ln n}{4n} \quad \text{w.p.1}$$

for all sufficiently large n and any $\varepsilon > 0$. Using $\tilde{\Delta} = \Delta + O(\frac{1}{n})$, (22) can be further approximated by

$$\ln w_n(\Delta) = -n[H(\tfrac{1}{2}) - H(\tfrac{1}{2} \pm \Delta)] + \tfrac{1}{2}\ln\tfrac{n}{2\pi} + O(1)$$

A Taylor series expansion around $\Delta=0$ yields

$$H(\tfrac{1}{2}) - H(\tfrac{1}{2}\pm\Delta) \;=\; 2\Delta^2 + O(\Delta^4) \;\leq\; (1+\varepsilon)\tfrac{\ln\ln n}{2n} + O((\tfrac{\ln\ln n}{4n})^2)$$

which implies

$$\ln w_n(\Delta_n) - \tfrac{1}{2}\ln\tfrac{n}{2\pi} + \ln\ln n \;\geq\; \tfrac{1}{2}(1-\varepsilon)\ln\ln n + O(1) \;\longrightarrow\; \infty \quad \text{for} \quad \varepsilon < 1$$

which implies (i) by exponentiation.

(ii) (a) Convexity of $\ln\Gamma(x)$ implies that $\ln w_n(\Delta)$ is concave and symmetric in Δ, hence $\ln w_n(\Delta)$ assumes its global maximum at $\Delta=0$. (b) From (22) it follows that $\ln w_n(0) = \frac{1}{2}\ln\frac{n}{2\pi} + (2\alpha-1)H(\frac{1}{2}) + C_\alpha + O(\frac{1}{n})$. (c) $(2n\Delta_n)_{n=1}^\infty$ is a symmetric random walk, hence infinitely often passes zero w.p.1. (a) and (b) imply the $\leq$ and (b) and (c) the $\geq$ in $\limsup_n[\ln w_n(\Delta_n) - \frac{1}{2}\ln\frac{n}{2\pi}] = (2\alpha-1)\ln 2 + C_\alpha$ w.p.1. Exponentiation yields (ii).

(iii) Since $\Delta_n \to \dot q_0 - \frac{1}{2}$ w.p.1, (22) implies $2nc^2 + \ln w_n \sim n[2c^2 - H(\frac{1}{2}) + H(\dot q_0)] \to -\infty$, since $H(\frac{1}{2}) - H(\dot q_0) \geq 2(\dot q_0 - \frac{1}{2})^2 > 2c^2$. The asymptotic representation also holds for $n_0=0$ or $n_1=0$, hence (iii) follows by exponentiation for all $\dot q_0$, including 0 and 1. □

Asymptotics of the evidence $p(D)$. The typical use of the posterior $p(x|D)$ is as an estimate for the unknown true distribution $\dot q(x)$. This makes sense if $p(x|D)$ is close to $\dot q(x)$. We show that the finite tree mixture model is indeed consistent in the sense that $p(x|D)$ converges[2] to $\dot q(x)$ and the posterior of $q()$ concentrates around the true value $\dot q()$ for $n\to\infty$.

Theorem 4 (Evidence asymptotics) *For fixed $m<\infty$ and $n_z\to\infty$, the posterior $p_z(x|D_z)\to\dot q_z(x)$ for all $x\in\Gamma_z$. Furthermore, for the evidence w.p.1 we have*

$$p_z(D_z) \;\begin{cases} \stackrel{poly.}{\longrightarrow} & u \quad \textit{for uniform } \dot q_z() \textit{ and } l<m, \\ \equiv & 1 \quad \textit{for } l=m, \\ \stackrel{exp.}{\longrightarrow} & \infty \quad \textit{for non-uniform } \dot q_z() \textit{ provided } s>0. \end{cases}$$

Proof by induction on l. We have to show slightly more, namely also that $p_z(D_z,x) \to c \in \{u,1,\infty\}$. For $l=m$, the theorem is obvious, since $\dot q_z(x)$ must be uniform on Γ_z and $p_z(D_z)\equiv 1 \equiv p_z(D_z,x)$, hence $p_z(x|D_z)\equiv 1 \equiv \dot q_z(x)$. Now assume the theorem holds for Γ_{z0} and Γ_{z1} and $l<m$. We show that it then also holds for Γ_z. Assume $u>0$ first.

(a) Assume first, that $\dot q_z()$ is uniform. This implies that also $\dot q_{z0}()$ and $\dot q_{z1}()$ are uniform, hence $n_{z0},n_{z1}\to\infty$, hence by induction hypothesis, $p_{z0}(D_{z0})$ and $p_{z1}(D_{z1})$ are bounded. Further, $w_{n_z}(\Delta_z) \stackrel{poly.}{\longrightarrow} \infty$ for $n_z\to\infty$ (by Theorem 3(i) and (ii)).

[2] All $\stackrel{n\to\infty}{\longrightarrow}$ statements hold with probability 1 (w.p.1).

Hence, $p_z(D_z) \stackrel{poly.}{\longrightarrow} u$ from (13). Similarly $p_z(D_z,x) \to u$, hence $p_z(x|D_z) \to 1 \equiv \dot{q}_z(x)$ for $x \in \Gamma_z$.

(b) We now consider the case of non-uniform $\dot{q}_z()$. (i) Consider the case that $\dot{q}_{z0}()$ or $\dot{q}_{z1}()$ (or both) are non-uniform first. $p_{z0}(D_{z0}) \geq u > 0$ and $p_{z1}(D_{z1}) \geq u > 0$, and one of them diverges exponentially. Since w_{n_z} grows at most with $O(\sqrt{n_z})$ by Theorem 3(ii), we see from (13) that also $p_z(D_z) \sim s \cdot p_{z0}(D_{z0}) p_{z1}(D_{z1}) / w_{n_z}$ diverges exponentially, and similarly $p_z(D_z,x)$. (ii) If both $\dot{q}_{z0}()$ and $\dot{q}_{z1}()$ are uniform, then $q_{z0} \neq \frac{1}{2}$, since we assumed non-uniform $\dot{q}_z()$. This implies bounded $p_{z0}(D_{z0})$ and $p_{z1}(D_{z1})$, but exponentially diverging w_n^{-1} by Theorem 3(iii). Hence, again, $p_z(D_z)$ and similarly $p_z(D_z,x)$ diverge exponentially. In both cases, (i) and (ii), assuming w.l.g. $x \in \Gamma_{z0}$, the ratio is

$$p_z(x|D) \sim \tfrac{w_{n_z}}{w_{n_z+1}} \cdot p_{z0}(x|D_{z0}) = 2 \cdot \tfrac{n_{z0}+\alpha}{n_z+2\alpha} \cdot p_{z0}(x|D_{z0}) \sim 2 \cdot \dot{q}_{z0} \cdot \dot{q}_{z0}(x) = \dot{q}_z(x) \qquad (23)$$

See (5) for notation and how the density factor 2 disappears. For $u = 0$, (23) holds for any $\dot{q}_z()$. Further, $p_z(D_z) \stackrel{exp.}{\longrightarrow} \infty$ still holds, since $p_{z0}(D_{z0})$ tends not faster than polynomially to zero by induction. □

Theorem 5 (Posterior consistency) *The posterior of $\vec{q}_{z*}$ concentrates for $n_z \to \infty$ around the true value $\dot{\vec{q}}_{z*}$ w.p.1., i.e.*[3]

$$p_z(\vec{q}_{z*}|D_z) \quad \stackrel{n_z \to \infty}{\underset{w.p.1}{\longrightarrow}} \prod_{y \in \mathbb{B}_1^{m-l}} \delta(q_{zy} - \dot{q}_{zy}) \qquad (24)$$

Proof. We prove consistency (24) by induction over l. For $\ell(z) = m$ the l.h.s. and r.h.s. are formally 1, since a density over an empty space and an empty product are defined as 1. Assume that consistency (24) holds for $\ell(z0) = \ell(z1) = l+1$. For $n_z \to \infty$, the Beta concentrates around $\frac{n_{z0}}{n_z} \to \dot{q}_{z0}$ and $\frac{n_{z1}}{n_z} \to \dot{q}_{z1}$ w.p.1:

$$\text{Beta}(q_{z0}, q_{z1}|n_{z0}+\alpha, n_{z1}+\alpha) \to \delta(q_{z0} - \dot{q}_{z0}) \delta(q_{z1} - \dot{q}_{z1})$$

Inserting this and (24) for $z0$ and $z1$ into the r.h.s. of recursion (15) we get

$$p_z(\vec{q}_{z*}|D_z) \to \frac{u}{p_z(D_z)} \prod_{y \in \mathbb{B}_1^{m-l}} \delta(q_{zy} - \tfrac{1}{2}) + \Big(1 - \frac{u}{p_z(D_z)}\Big) \prod_{y \in \mathbb{B}_1^{m-l}} \delta(q_{zy} - \dot{q}_{zy}) \qquad (25)$$

For uniform $\dot{\vec{q}}_{z*}$, i.e. $\frac{1}{2} = \dot{q}_{zy}$ $\forall y \in \mathbb{B}_1^{m-l}$ the r.h.s. reduces to the r.h.s. of (24). For non-uniform $\dot{\vec{q}}_{z*}$, Theorem 4(iii) shows that $p_z(D_z) \to \infty$ (exponentially), and the r.h.s. of (25) converges (rapidly) to the r.h.s. of (24). □

[3] The topology of weak convergence or convergence in distribution is used.

5 More Quantities of Interest

In this section we introduce further quantities of interest and present recursions for them. They all can be written as expectations

$$E_z[f(\vec{q}_{z*})|D_z] \;:=\; \int f(\vec{q}_{z*})p_z(\vec{q}_{z*}|D_z)d\vec{q}_{z*} \tag{26}$$

for suitable functions f (and similarly for $P_z[...|...]=P[...|\Gamma_z...]$). For instance for the evidence (13) we used $f(\vec{q}_{z*})\equiv 1$ and for the posterior (15) formally $f(\vec{q}_{z*})=\prod_y \delta(q_{zy}-q'_{zy})$. Below we consider the model dimension, cell number, tree height, cell size, and moments. The derivations of the recursion all follow the same scheme, inserting the (recursive) definition of f and the recursion (15) into the r.h.s. of (26), and rearranging and identifying terms. These details will be omitted.

Model dimension and cell number. As discussed in Section 2, the effective dimension of $\vec{q}_*$ is the number of components that are not forced to $\frac{1}{2}$ by (6). Note that a component may be "accidentally" $\frac{1}{2}$ in (8), but since this is an event of probability 0, we don't have to care about this subtlety. So the effective dimension $N_{\vec{q}_{z*}}=\#\{q\in\vec{q}_{z*}:q\neq\frac{1}{2}\}$ of $\vec{q}_{z*}$ can be given recursively as

$$N_{\vec{q}_{z*}} \;=\; \begin{cases} 0 & \text{if} \quad \ell(z)=m \quad \text{or} \quad q_{z0}=\frac{1}{2} \\ 1+N_{\vec{q}_{z0*}}+N_{\vec{q}_{z1*}} & \text{else} \end{cases} \tag{27}$$

The effective dimension is zero if $q_{z0}=\frac{1}{2}$, since this implies that the whole tree Γ_z has $q_{zy}=\frac{1}{2}$ due to (12). If $q_z\neq\frac{1}{2}$, we add the effective dimensions of subtrees Γ_{z0} and Γ_{z1} to the root degree of freedom $q_{z0}=1-q_{z1}$. We may be interested in the expected effective dimension $E[N_{\vec{q}_*}|D]$. Inserting (27) ($f(\vec{q}_{z*})=N_{\vec{q}_{z*}}$) and (15) into the r.h.s. of (26) we can prove the following recursion for the expected effective dimension

$$E_z[N_{\vec{q}_{z*}}|D_z]=g_z(D_z)[1+E_{z0}[N_{\vec{q}_{z0*}}|D_{z0}]+E_{z1}[N_{\vec{q}_{z1*}}|D_{z1}]] \tag{28}$$

Read: The expected dimension of $\vec{q}_{z*}$ (l.h.s.) equals to 1 for the root degree of freedom plus the expected dimensions in the left and right subtrees, multiplied with the probability $g_z(D_z)$ of splitting Γ_z (r.h.s.). The recursion terminates with $E_z[N_{\vec{q}_{z*}}|D_z]=0$ when $\ell(z)=m$. Higher (central) moments like the variance can be computed similarly. One can also compute the whole distribution $(P[N_{\vec{q}_*}=k|D])_{k\in\mathbb{N}}$ by convolution. Inserting

$$f(\vec{q}_{z*}) \;=\; \delta_{N_{\vec{q}_{z*}},k+1} \;=\; \begin{cases} 0 \quad \text{if} \quad l=m \quad \text{or} \quad q_{z0}=\frac{1}{2}, \\ \sum_{i=0}^{k}\delta_{N_{\vec{q}_{z0*}},i}\delta_{N_{\vec{q}_{z1*}},k-i} \quad \text{else}, \end{cases} \tag{29}$$

and (15) into (26) we get

$$\begin{aligned} &P_z[N_{\vec{q}_{z*}}=0|D_z] \;=\; 1-g_z(D_z), \qquad\qquad \text{for} \quad l<m, \\ &P_z[N_{\vec{q}_{z*}}\underset{k}{=}k+1|D_z] \;= \\ &g_z(D_z)\cdot\sum_{i=0} P_{z0}[N_{\vec{q}_{z0*}}=i|D_{z0}]\cdot P_{z1}[N_{\vec{q}_{z1*}}=k-i|D_{z1}], \\ &P_z[N_{\vec{q}_{z*}}=k|D_z]=\delta_{k0}:=\left\{\begin{smallmatrix}1 \text{ if } k=0\\ 0 \text{ if } k>0\end{smallmatrix}\right\} \quad \text{for} \quad l=m. \end{aligned} \tag{30}$$

Read: The probability that tree Γ_z has dimension $k+1$ equals the probability of splitting, times the probability that left subtree has dimension i, times the probability that right subtree has dimension $k-i$, summed over all possible i. Again, this follows from inserting (29) and (15) into (26).

Let us define a cell or bin as a maximal volume on which $q()$ is constant. Then the model dimension is 1 less than the number of bins (due to the probability constraint). Hence adding 1 to the above quantities we also have expressions for the expected number of cells and distribution.

Tree height and cell size. The effective height of tree $\vec{q}_{z*}$ at $x \in \Gamma_z$ is also an interesting property. If $q_{z0} = \frac{1}{2}$ or $\ell(z) = m$, then the height $h_{\vec{q}_{z*}}(x)$ of tree $\vec{q}_{z*}$ at x is obviously zero. If $q_{z0} \neq \frac{1}{2}$, we take the height of the subtree $\vec{q}_{zx_{l+1}*}$ that contains x and add 1:

$$h_{\vec{q}_{z*}}(x) \;=\; \begin{cases} 0 & \text{if} \quad \ell(z) = m \quad \text{or} \quad q_{z0} = \frac{1}{2} \\ 1 + h_{\vec{q}_{zx_{l+1}*}}(x) & \text{else} \end{cases}$$

One can show that the tree height at $x \in \Gamma_z$ averaged over all trees $\vec{q}_{z*}$ is

$$E_z[h_{\vec{q}_{z*}}(x)|D_z] \;=\; g_z(D_z)\Big[1 + E_{zx_{l+1}}[h_{\vec{q}_{zx_{l+1}*}}(x)|D_{zx_{l+1}}]\Big]$$

We may also want to compute the tree height averaged over all $x \in \Gamma_z$.

$$\bar{h}_{\vec{q}_{z*}} \;:=\; \int h_{\vec{q}_{z*}}(x) q(x|\Gamma_z) dx \;=\; \begin{cases} 0 & \text{if} \quad \ell(z) = m \quad \text{or} \quad q_{z0} = \frac{1}{2} \\ 1 + q_{z0} \cdot \bar{h}_{\vec{q}_{z0*}} + q_{z1} \cdot \bar{h}_{\vec{q}_{z1*}} & \text{else} \end{cases}$$

$$\begin{aligned} E_z[\bar{h}_{\vec{q}_{z*}}|D_z] \;=\; g_z(D_z)\Big[1 &+ \frac{n_{z0}+\alpha}{n_z+2\alpha} E_{z0}[\bar{h}_{\vec{q}_{z0*}}|D_{z0}] \\ &+ \frac{n_{z1}+\alpha}{n_z+2\alpha} E_{z1}[\bar{h}_{\vec{q}_{z1*}}|D_{z1}]\Big] \end{aligned}$$

with obvious interpretation: The expected height of a subtree is weighted by its relative importance, that is (an estimate of) its probability. The recursion terminates with $E_z[h_{\vec{q}_{z*}}|D_z] = 0$ when $\ell(z) = m$. We can also compute intra and inter tree height variances.

Finally consider the average cell size or volume v. Maybe more useful is to consider the logarithm $-\log_2|\Gamma_z| = \ell(z)$, since otherwise small volumes can get swamped in the expectation by a single large one. Log-volume $v_{\vec{q}_{z*}} = \ell(z)$ if $\ell(z) = m$ or $q_{z0} = \frac{1}{2}$, and else recursively $v_{\vec{q}_{z*}} = q_{z0} v_{\vec{q}_{z0*}} + q_{z1} v_{\vec{q}_{z1*}}$. We can reduce this to the tree height, since $v_{\vec{q}_{z*}} = \bar{h}_{\vec{q}_{z*}} + \ell(z)$, in particular $v_{\vec{q}_*} = \bar{h}_{\vec{q}_*}$

Moments in x. Assume data $x \in \Gamma = [0,1)$ are sampled from $q()$. The mean and variance of x w.r.t. $q()$ are important statistical quantities. More generally consider

$$f(\vec{q}_{z*}) \;=\; M_{\vec{q}_{z*}} \;:=\; \frac{1}{|\Gamma_z|} \int_{\Gamma_z} M(x) q_z(x) dx$$

($\frac{1}{|\Gamma_z|}\int q_z(x)dx=1$). Since $q()$ is itself random, it is natural to consider the p-expected q-mean (26) of M

$$\mathbb{E}_z[M(x)|D_z] \;:=\; E_z[M_{\vec{q}_{z*}}|D_z] \;=\; \frac{1}{|\Gamma_z|}\int_{\Gamma_z} M(x)p_z(x|D_z)dx \tag{31}$$

Inserting recursion (13) for $p_z(D_z,x)$ using (4) in (31) we get

$$\mathbb{E}_z[M(x)|D_z]=\frac{u\cdot\bar{M}_z}{p_z(D_z)}+g_z(D_z)\Big[\frac{n_{z0}+\alpha}{n_z+2\alpha}\mathbb{E}_{z0}[M(x)|D_{z0}]+\frac{n_{z1}+\alpha}{n_z+2\alpha}\mathbb{E}_{z1}[M(x)|D_{z1}]\Big], \tag{32}$$

again with obvious interpretation: The expectation of M is a mixture of a uniform expectation $\bar{M}_z:=\frac{1}{|\Gamma_z|}\int_{\Gamma_z}M(x)dx$ and the weighted average of expectations in left and right subtree. The recursion terminates with $\mathbb{E}_z[M(x)|D_z]=\bar{M}_z$ when $l=m$. Examples: For $M(x)\equiv 1$, both sides of (32) evaluate to 1 as it should. For the kth moment of x, $M(x)=x^k$ we have $\bar{M}_z=[(z+1)^{k+1}-z^{k+1}]/[2^{kl}(k+1)]$, where $z=2^l0.z$ is interpreted as an integer in binary representation. The distribution function $P_z[x\leq a|D_z]$ is obtained for $M(x)=\{^{1 \text{ if } x\leq a}_{0 \text{ if } x>a}$. For $a\in\Gamma_z$ we have $\bar{M}_z=2^la-z$. Since $M(x)$ is constant on $\Gamma_z\not\ni a$, we have $\mathbb{E}_z[M(x)|D_z]=\{^{1 \text{ if } a\geq 0.z+2^{-l}}_{0 \text{ if } a<0.z}$ in this case, hence only one recursion in (32) needs to be expanded (since $a\notin\Gamma_{z0}$ or $a\notin\Gamma_{z1}$).

6 Infinite Trees ($m\to\infty$)

Motivation. We have chosen an (arbitrary) finite tree height m in our setup, needed to have a well-defined recursion start at the leaves of the trees. What we are really interested in are infinite trees ($m=\infty$). Why not feel lucky with finite m? First, for continuous domain Γ (e.g. interval [0,1)), our tree model contains only piecewise constant models. The true distribution $\dot{q}()$ is typically non-constant and continuous (Beta, normal, ...). Such distributions are outside a finite tree model class (but inside the infinite model), and the posterior $p(x|D)$ cannot converge to the true distribution, since it is also piecewise constant. Hence all other estimators based on the posterior are also not consistent. Second, a finite m violates scale invariance (a non-informative prior on Γ_z should be the same for all z, apart from scaling). Finally, having to choose the "right" m may be worrisome.

For increasing m, the cells Γ_x become smaller and will (normally) eventually contain either only a single data item, or be empty. It should not matter whether we further subdivide empty or singleton cells. So we expect inferences to be independent of m for sufficiently large m, or at least the limit $m\to\infty$ to exist. In this section we show that this is essentially true, but with interesting exceptions.

Prior inferences ($D=\emptyset$). We first consider the prior (zero data) case $D=\emptyset$. Recall that $z\in\mathbb{B}_0^m$ is some node and $x\in\mathbb{B}^m$ a leaf node. Normalization implies $p_z(\emptyset)=1$ for all z, which is independent of m, hence the prior evidence exists for $m\to\infty$ (see below for a formal proof). This is nice, but hardly surprising.

The prior effective model dimension $N_{\vec{q}*}$ is more interesting. $D=\emptyset$ implies $D_z=\emptyset$ implies $n_z=0$ implies $w(n_{z0},n_{z1})=1$ implies a chance $g_z(\emptyset)=1-u=s$ for a split (see (16)). The recursion (28) reduces to

$$E_z[N_{\vec{q}_{z*}}] = s[1 + E_{z0}[N_{\vec{q}_{z0*}}] + E_{z1}[N_{\vec{q}_{z1*}}]]$$

with $E_z[N_{\vec{q}_{z*}}]=0$ for $l\equiv\ell(z)=m$ and is easily solved:

$$E_z[N_{\vec{q}_{z*}}] \;=\; s(1+2s(1+...)) \;=\; s\frac{1-(2s)^{m-l}}{1-2s} \;=\; \begin{cases} \overset{m\to\infty}{\longrightarrow} \frac{s}{1-2s} & \text{if} \quad s<\frac{1}{2} \\ \frac{1}{2}(m-l) \overset{lin.}{\to} \infty & \text{if} \quad s=\frac{1}{2} \\ \overset{exp.}{\to} \infty & \text{if} \quad s>\frac{1}{2} \end{cases}$$

This can be understood as follows. Consider trees truncated at nodes with uniform distribution. Assume that there are $k(l)$ nodes at height l. With probability u the node is a leaf, and with probability s it has two children. So the expected number of nodes at height $l+1$ is $k(l+1)=2s\cdot k(l)$. So the number of nodes exponentially increases/decreases with l for $s>\frac{1}{2}/s<\frac{1}{2}$, which results in a infinite/finite total number of nodes (=dimension). The linear divergence for $s=\frac{1}{2}$ looks alerting (overfitting), but we now show that the distribution exists and an infinite expectation is actually a good sign. Recursion (30) reads

$$P_z[N_{\vec{q}_{z*}}=k+1] = s\sum_{i=0}^{k} P_{z0}[N_{\vec{q}_{z0*}}=i]\cdot P_{z1}[N_{\vec{q}_{z1*}}=k-i]$$

with $P_z[N_{\vec{q}_{z*}}=0]=u$ for $l<m$ and $P_z[N_{\vec{q}_{z*}}=k]=\delta_{k0}$ for $l=m$. So the recursion terminates in recursion depth $\min\{k+1,m-l\}$. Hence $P_z[N_{\vec{q}_{z*}}=k+1]$ is the same for all $m>l+k$, which implies that the limit $m\to\infty$ exists. Furthermore, recursion and termination are independent of z, hence also $a_k:=P_z[N_{\vec{q}_{z*}}=k]$. So we have to solve the recursion

$$a_{k+1} = s\sum_{i=0}^{k} a_i\cdot a_{k-i} \quad \text{with} \quad a_0=u \tag{33}$$

The first few coefficients can be bootstrapped by hand, e.g. for $s=\frac{1}{2}$ we get ($\vec{a}=\frac{1}{2},\frac{1}{8},\frac{1}{16},\frac{5}{128},\frac{7}{256},\frac{21}{1024},\frac{33}{2048},...$). A closed form can also be obtained: Inserting (33) into $f(x):=\sum_{k=0}^{\infty}a_kx^{k+1}$ we get $f(x)=ux+sf^2(x)$ with solution $f(x)=\frac{1}{2s}[1-\sqrt{1-4sux}]$, which has Taylor expansion coefficients

$$a_k = 2u(-4su)^k\binom{1/2}{k+1} = \frac{u(4su)^k}{(k+1)4^k}\binom{2k}{k} \sim \frac{u(4su)^k}{\sqrt{\pi}}\,k^{-3/2}$$

For $s\leq\frac{1}{2}$, $(a_k)_{k\in\mathbb{N}_0}$ is a well-behaved properly normalized probability measure ($\sum_k a_k=f(1)=1<\infty$). For $s<\frac{1}{2}$ it decreases exponentially in k, implying that all moments exist and indicating strong bias towards simple models. For $s=\frac{1}{2}$, $a_k\sim k^{-3/2}$ decreases polynomially in k, too slow for the expectation $E[N_{\vec{q}*}]=\sum_k k\cdot a_k=\infty$ to exist, but this is exactly how a proper non-informative prior on $\mathbb{N}$ should look like: as

uniform as possible, i.e. slowly decreasing. Further, $P[N_{\vec{q}_*}<\infty]=\sum_k a_k=1$ shows that the effective dimension is almost surely finite, i.e. infinite (Polya) trees have probability zero for $s\leq\frac{1}{2}$. On the other hand, for $s>\frac{1}{2}$, we have $P[N_{\vec{q}_*}=\infty]=1-f(1)=2-\frac{1}{s}$, i.e. a non-zero probability for infinite trees. The reason why a_k also exponentially decreases in this case is that as deeper a tree grows as less likely it stays finite. These results are consistent with the expected model dimension. They indicate a proper behavior of our model for $s\leq\frac{1}{2}$ and in particular for $s=\frac{1}{2}$.

For the tree height we have $E_z[h_{\vec{q}_{z*}}(x)]=0$ if $l=m$ and otherwise

$$\begin{aligned} E_z[h_{\vec{q}_{z*}}(x)] &= s[1+E_{zx_{l+1}}[h_{\vec{q}_{zx_{l+1}*}}(x)]] \\ &= s+s^2+...+s^{m-l} \\ &= \begin{cases} s\cdot\frac{1-s^{m-l}}{1-s} \overset{m\to\infty}{\longrightarrow} \frac{s}{1-s} & \text{if} \quad s<1, \\ m-l \overset{m\to\infty}{\longrightarrow} \infty & \text{if} \quad s=1, \end{cases} \end{aligned}$$

i.e. the prior expected height is large/small if the splitting probability s is large/small. The same holds for the expected average height $E_z[\bar{h}_{\vec{q}_{z*}}]\to\frac{s}{1-s}$. This is the first case where the result is not independent of m for large finite m, but it converges for $m\to\infty$, what is enough for our purpose. Note that a finite expected tree height even for $1>s\geq\frac{1}{2}$ is consistent with an infinite model dimension, since $h=\infty$ only for a vanishing fraction of tree branches $x\in\Gamma$, i.e. for a set of measure zero.

The prior moments $M(x)$ are easy to compute: Since $p_z(x)\equiv 1$, we get $\mathbb{E}_z[M(x)]=\bar{M}_z$.

Multi-points $D=(x^1,...,x^1)$. The next situation we analyze is multiple points $D=(x^1,...,x^1)$, where all data points are identical. For continuous spaces and non-singular priors, the probability of such an event is zero, so this scenario does not seem particulary interesting, but: When computing posteriors, x is not chosen randomly but deliberately, so in computing $p(D,x)$, x could be equal to x^1 (although again only with probability zero). When computing higher moments we need $p(D,x,...,x)$ and definitely encounter multiple points. Also, multi-points help to understand the case when x or x^i comes very close to x^1. Also the true prior may be singular causing $x^1=x^2$ with non-zero probability. Finally, the multi-point case includes $n=1$, which we have to analyze in any case.

D_z is either $\emptyset$ or D. $D_z=\emptyset$ has been studied in the last §, so we assume $D_z=D$. Either $D_{z0}=\emptyset$ or $D_{z1}=\emptyset$. W.l.g. we assume the latter. Then $D_{z0}=D$, $n_{z0}=n$, $n_{z1}=0$, which implies $w(n_{z0},n_{z1})=2^{-n}\frac{\Gamma(n+2\alpha)\Gamma(\alpha)}{\Gamma(n+\alpha)\Gamma(2\alpha)}$. Defining $\bar{w}:=s/w$, recursion (13) reduces to

$$p_z(D) \;=\; u+\bar{w}\cdot p_{z0}(D) \;=...=\; u\frac{1-\bar{w}^{m-l}}{1-\bar{w}}+\bar{w}^{m-l}$$

$$\begin{cases} \equiv 1 & \text{if} \quad \bar{w}=s \\ \to\frac{u}{1-\bar{w}} & \text{if} \quad \bar{w}<1 \\ =u(m-l)+1\overset{lin.}{\longrightarrow}\infty & \text{if} \quad \bar{w}=1 \\ \to\frac{\bar{w}-s}{\bar{w}-1}\bar{w}^{m-l}\overset{exp.}{\longrightarrow}\infty & \text{if} \quad \bar{w}>1 \end{cases}$$

So the evidence exists for $m\to\infty$ iff $\bar w<1$. For $n=0$ and $n=1$ we have $\bar w=s<1$ (excluding $s=1$), hence $p(\emptyset)\equiv 1$ is correctly normalized, as claimed in the previous §, and $p(x)\equiv 1$ is uniform as symmetry demands. For double points $n=2$ the evidence $p(x,x)$ is still finite iff $\bar w=s\cdot\frac{\alpha+1}{\alpha+1/2}<1$. The latter is true for $s\leq\frac12$ and all $\alpha>0$. For any n but $\alpha=1$, $\bar w<1$ iff $s<(n+1)2^{-n}$, i.e. higher multi-point evidences only exist for exponentially small s. If we want $\bar w<1\ \forall\alpha>0$, $s<2^{1-n}$ has to be even smaller. This follows from $\bar w$ being increasing in n and decreasing in α ($\bar w\searrow s$ for $\alpha\to\infty$). To conclude: For every fixed α and s, multiplicity n of points must not be too high, but for any n one can choose α sufficiently large or s sufficiently small so that the n-multi-point evidence and hence the n^{th} moments exist.

For the mean of $M(x)$ we get

$$\mathbb{E}_z[M(x)|D] \;=\; (1-\tilde w)\bar M_z + \tilde w\tfrac{\alpha}{n+2\alpha}\bar M_{z0} + (\tilde w\tfrac{n+\alpha}{n+2\alpha})\mathbb{E}_{z1}[M(x)|D]$$

if $x^1\in\Gamma_{z1}$ and similarly for $x^1\in\Gamma_{z0}$, where $\tilde w=\min\{\bar w,1\}$. This is not a self-consistency equation, but since $\tilde w\frac{n+\alpha}{n+2\alpha}<1$, the linear recursion converges exponentially to the exact value.

Multi-points for $\alpha=1$ and $s=\frac12$. We present some more results for the most interesting case $\alpha=1$ and $s=\frac12$, which we will also investigate numerically.

We see that $p(\emptyset)\equiv 1$ is correctly normalized, and $p(x)\equiv 1$ is uniform as symmetry demands. For double points, the evidence $p(x,x)\to\frac32$ is still finite. It diverges linearly for triple points and exponentially for quadruple-and-higher points. So $q(x)$ has finite prior mean $E[q(x)]=1$ and variance $\mathrm{Var}[q(x)]=\frac12$ (Section 2). The skewness and kurtosis are infinite, indicating a heavy tail, as desired for a non-informative prior.

Since $p(D)\equiv 1$ and $w=1$ for $n=1$ are the same as for the $n=0$ case, all prior $n=0$, $m\to\infty$ results remain valid for $n=1$: $g(x)=\frac12$, $E[N_{\vec q_*}|x]=\infty$, $P[N_{\vec q_*}=k|x]=a_k$, and $E[h_{\vec q_*}(x)|x]\to 1$.

For $n=2$ we get $g(x,x)=\frac23$, $b_k:=P[N_{\vec q_*}=k|x,x]$, $b_0=1-g(x,x)=\frac13$, $b_{k+1}=\frac23\sum_{i=0}^k b_i a_{k-i}=(\frac19,\frac{7}{108},\frac{29}{648},...)$, $h(x):=\sum_{k=0}^\infty b_k x^{k+1}=\frac13[1+2h(x)f(x)]=\frac{x}{1+2\sqrt{1-x}}$, and $E[h_{\vec q_*}(x)|x,x]\to 2$.

For $n\geq 3$ we have $g_z(D)=1$, $b_0=0\Rightarrow b_k\equiv 0\Rightarrow P[N_{\vec q_*}<\infty|D]=0$, $E[h_{\vec q_*}(x)|D]=m\to\infty$ for $x=x^1$. The tree at x has infinite height and singular distribution.

For all n we have $p(D)\geq 1$, $g(D)\geq\frac12$, $E[N_{\vec q_*}|D]\sim 2g(D)$ for $m\to\infty$,

Not a double point, but also straightforward to compute is $p(x,y)=\frac32-(\frac23)^{l+1}$ if $x\in\Gamma_{z0}$ and $y\in\Gamma_{z1}$, i.e. x and y separate at level $l=\ell(z)$, consistent with $p(x,x)=\frac32$ ($l\hat=\infty$).

General D. We now consider general D. In order to compute $p(D)$ and other quantities, we recurse (13) down the tree until D_z is a multi-point $D_z=D_z^=:=(x',...,x')$ with $x'\in\Gamma_z$. We call the depth $m_{x'}:=\ell(z)$ at which this happens, the separation level. If we consider $n_z\in\mathbb{N}_0$, this also includes the most important empty

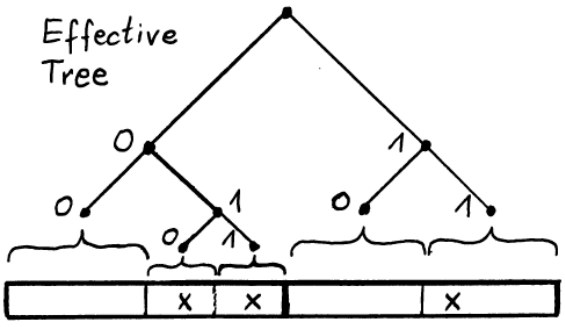

Figure 4: Effectively recursed tree. Closed form solutions are used for intervals containing no or only a single (multi)point.

and singleton case. In this way, the recursion always terminates. For instance, for $\Gamma = [0,1)$, if $\varepsilon := \min\{|x^i - x^j| : x^i \neq x^j \text{ with } x^i,x^j \in D\}$ is the shortest distance, then $m_{x'} < \log_2 \frac{2}{\varepsilon} =: m_0 < \infty$, since $\varepsilon > 0$. At the separation level we can insert the derived formulas (for evidence, posterior, dimension, height, moments) for multi-points (Figure 4). Note, there is no approximation here. The procedure is exact, since we analytically computed the infinite recursion for multi-points.

So we have devised a finite procedure for exactly computing all quantities of interest. In the worst case, we have to recurse down to level m_0 for each data point, hence our procedure has computational complexity $O(n \cdot m_0)$. For non-singular prior, the time is actually $O(n)$ with probability 1. So, inference in our mixture tree model is *very* fast. Polya trees for suitable Beta prior (should) admit similar algorithms.

Multi-point divergences. Consider again $s = \frac{1}{2}$ and $\alpha = 1$. Since $E_z[N_{\vec{q}_{z*}}|D_z^=] = \infty$ recurses up, we have $E[N_{\vec{q}_*}|D] = \infty$ for all D. We now discuss possible divergences caused by true multi-points $n_z^= > 2$ in D. $E_z[h_{\vec{q}_{z*}}(x)|D_z^=] = \infty$ recurses up to $E[h_{\vec{q}_*}(x)|D] = \infty$. Similarly, $P_z[N_{\vec{q}_{z*}} < \infty|D_z^=] = 0$, recurses up to $P[N_{\vec{q}_*} < \infty|D] = 0$, since $g_z(D_z) = 1$ along the path.

There are interesting cases where $p(D) = \infty$, but posteriors are finite, since infinities cancel out. These are $p(x|D)$ if x occurs at most once in D, and $p(x,x,|D)$ if $x \notin D$, otherwise $p(x,...,x|D) = \infty$. This is is very welcome: $E[q(x)|D] < \infty$ for all x, if D contains only singletons, which is true w.p.1 for all regular $\dot{q}()$. The posterior variance of $q(x)$ is finite iff $x \notin D$, which could be better. We adapt recursion (13) by scaling $p(D)$, $p(D,x)$ and $p(D,x,x)$ with the same constant $c_m \to 0$ such that they stay finite, which works since D, (D,x) and (D,x,x) have the same triple, quadruple, ... points. Choose l large enough (separation level) so that D_z is a triple-or-higher point, and w.l.g. D_{z1} is empty. Then the recursion (13) reduces to recursion (34). The $\frac{1}{2}+$ in (13) gets swamped by $p_{z0}(D_{z0}^=) = \infty$ and can be dropped. The remaining recursion is just a multiplication with $\bar{w}$, which allows us to rescale $p_z(D_z^=)$ to $p_z(D_z^=) = \bar{w}^{-l}$ (cf. (34)). We return a flag in the recursion indicating that the $\frac{1}{2}+$ shall be dropped along the path back to the root, since the true original $p_z(D_z^=)$ was infinite and would have swamped them all. We compute $p(D)$, $p(D,x)$, and

$p(D,x,x)$ in this way to arrive at finite posteriors (not forgetting that the numbers we get for the evidences themselves are fictitious).

A heuristic way of regularizing our model to yield always finite results could be to cut the recursion short at the separation level by definition, and then assign some regular (e.g. uniform) distribution to this leaf. Since the most important quantities are finite anyway, we refrain from such a data-dependent non-Bayesian hack. Better is to assign a smaller prior weight $s<\frac{1}{2}$ to a split; then more (higher) moments become finite.

Consistency ($n\to\infty$). What remains to be shown is posterior consistency for $m=\infty$ similarly to the $m<\infty$ case. We will show weak consistency in the sense that $p(\{q_z\}|D)$ concentrates around $\{\dot q_z\}$, where $\{q_z\}$ is a *finite* collection of branching probabilities. Consistency holds because the recursion for $p(\{q_z\}|D)$ terminates at a depth independent of m (for sufficiently large m), so we are effectively in the finite tree case. The only difficulty is that the recursion involves $p_z(D_z)$ which still has recursion depth (i.e. depends on) m. One solution could be to assume that all observations have some finite precision $2^{-m'}$ and $x\in D\rightsquigarrow\Gamma_{x'}\in D'$, where $\ell(x')=m'$ and $x\in\Gamma_{x'}$. This would make all involved recursions terminate at depth m' and hence all recursions and results for finite m apply (with m' instead of m). More interesting is to keep D and treat the $m=\infty$ case properly. We show that for $n\to\infty$, the evidence and the posterior converge uniformly in $m<\infty$, which implies convergence also for $m=\infty$. We sometimes indicate the m-dependence of p_z by p_z^m and define $p_z^\infty:=\lim_{m\to\infty}p_z^m$ if the limit exists (possibly infinite), but mostly drop the superscript $m\le\infty$.

Theorem 6 (Weak consistency for infinite trees) *Let $\ell(z)=l\le m'<m+1\le\infty$ and $\vec q_{z*'}=(q_{zy}:y\in\mathbb{B}_1^{m'-l})$. Then evidence and p(oste)rior $p_z(...)$ exist for $m=\infty$ and have the following properties for $m\le\infty$, where convergence $\overset{n_z\to\infty}{\longrightarrow}$ holds w.p.1 and is uniform in m:*

i) The marginal prior $p_z(\vec q_{z*'})$ is independent of m,

ii) $p_z(D_z)\overset{n_z\to\infty}{\underset{exp.}{\longrightarrow}}\infty$ for $us>0$ and non-uniform $\dot q_z()$.

iii) $p_z(\vec q_{z*'}|D_z)\quad\overset{n_z\to\infty}{\underset{weak}{\longrightarrow}}\prod_{y\in\mathbb{B}_1^{m'-l}}\delta(q_{zy}-\dot q_{zy})$.

iv) $p_z(\Gamma_{y^1},...,\Gamma_{y^k}|D_z)\overset{n_z\to\infty}{\longrightarrow}\dot q_z(\Gamma_{y^1})\cdot...\cdot\dot q_z(\Gamma_{y^k})\quad$ for $y^i\in\mathbb{B}^*$ and $k\in\mathbb{N}$.

(iv) implies weak convergence of $p(x|D)$ to $\dot q(x)$ in the sense that $\int f(x)p(x|D)dx\to\int f(x)\dot q(x)dx$ for continuous functions f, by an argument similar to the proof of [Fab64, Thm.2.2]. Also the distribution function $P[x\le a|D]\to P[x\le a|\dot q]$.

Proof. We only have to consider $m<\infty$. The $m=\infty$ case follows for (i) by definition and for $(ii)-(iv)$ since convergence $\overset{n_z\to\infty}{\longrightarrow}$ is uniform in m.

(i) The prior marginal is

$$p_z(\vec{q}_{z*'}) \;=\; \int p_z(\vec{q}_{z*}) \prod_{y\in\mathbb{B}_1^{m-l}\setminus\mathbb{B}_1^{m'-l}} dq_{zy} \tag{35}$$

For $l=m'$ we have $p_z(\vec{q}_{z*'})=p_z(\emptyset)=1$ independent of m. For $l<m'$, inserting recursion (12) into (35) and using $\mathbb{B}_1^{m-l}\setminus\mathbb{B}_1^{m'-l}=\{0,1\}\times(\mathbb{B}_1^{m-(l+1)}\setminus\mathbb{B}_1^{m'-(l+1)})$ and (35) for $z0$ and $z1$ backwards, we get

$$p_z(\vec{q}_{z*'}) = u\cdot\prod_{y\in\mathbb{B}_1^{m'-l}}\delta(q_{zy}-\tfrac{1}{2}) + s\cdot\mathrm{Beta}(q_{z0},q_{z1}|\alpha,\alpha)p_{z0}(\vec{q}_{z0*'})p_{z1}(\vec{q}_{z1*'})$$

So the recursion of $p_z^m(\vec{q}_{z*'})$ and its termination is independent of m, hence $p_z^m(\vec{q}_{z*'})$ is independent of m, hence $p_z^\infty(\vec{q}_{z*'})$ exists and equals $p_z^m(\vec{q}_{z*'})$ for $m\geq m'$.

(ii) First note that $p_z(D_z)$ in general depends on m. Further, $p_z(D_z)\geq u>0\;\forall z$. Assume first that $\dot{q}_{z0}\neq\frac{1}{2}$. Then from (13) and Theorem 3(iii) we get

$$p_z(D_z) \;=\; u + s\cdot\frac{p_{z0}(D_{z0})p_{z1}(D_{z1})}{w_{n_z}(\Delta_z)} \;\geq\; \frac{su^2}{w_{n_z}} \overset{n_z\to\infty}{\underset{exp.}{\longrightarrow}} \infty \tag{36}$$

Divergence of $p_z(D_z)$ is uniform in m, since w_{n_z} is independent m. Now consider the more general case of non-uniform $\dot{q}_z()$, i.e. $\exists y:\dot{q}_{zy0}\neq\frac{1}{2}$. Then (36) implies $p_{zy}(D_{zy})\overset{exp.}{\longrightarrow}\infty$. Further, for any z, if $p_{z0}\overset{exp.}{\longrightarrow}\infty$, then $p_z\overset{exp.}{\longrightarrow}\infty$, since $p_{z1}\geq u>0$ and $w_{n_z}=O(\sqrt{n_z})$ by Theorem 3, and similarly if $p_{z1}\to\infty$. So by induction, $p_{zy}\overset{exp.}{\longrightarrow}\infty$ uniformly in m implies $p_z\overset{exp.}{\longrightarrow}\infty$ uniformly in m. Finally, $p_z^\infty(D_z)$ exists, since $p_z^m(D_z)$ is independent m beyond the data separation level, as shown earlier.

(iii) Similarly to the marginal prior recursion in (i) one can show that the recursion for the marginal posterior $p_z(\vec{q}_{z*'}|D_z)$ has the same form as the recursion (15) of $p_z(\vec{q}_{z*}|D_z)$ with m replaced by m'. Contrary to the prior, the posterior still depends on m through $p_z(D_z)$. Choosing m beyond the data separation level does not help since it increases with n_z. Nevertheless, the proof of (iii) is the same as for Theorem 5 with $m\rightsquigarrow m'$. Convergence is uniform in m, since convergence (25) and divergence of $p_z(D_z)$ are uniform. Finally, $p_z^\infty(\vec{q}_{z*'}|D_z)$ exists, since its recursion is finite (terminates at m') and $p_z^\infty(D_z)$ exists.

(iv) Choose $m'\geq\max\{\ell(y^1),...,\ell(y^k)\}$. Then

$$p_z(\Gamma_{y^1},...,\Gamma_{y^k}|D_z) \;=\; \int q_z(\Gamma_{y^1})\cdot...\cdot q_z(\Gamma_{y^k})\cdot p_z(\vec{q}_{z*'}|D_z)d\vec{q}_{z*'}$$

exists, and (iv) now follows from (iii) and $q_z(\Gamma_y)=q_{y_1}\cdot...\cdot q_{y_{1:\ell(y)}}$. □

7 The Algorithm

What it computes. In the last two sections we derived all necessary formulas for making inferences with our tree model. Collecting pieces together we get the

exact algorithm for infinite tree mixtures below. It computes the evidence $p(D)$, the expected tree height $E[h_{\vec{q}_*}(x)|D]$ at x, the average expected tree height $E[\bar{h}_{\vec{q}_*}|D]$, and the model dimension distribution $P[N_{\vec{q}_*}|D]$. It also returns the number of recursive function calls, i.e. the size of the explicitly generated tree. The size is proportional to n for regular distributions $\dot{q}$.

The BayesTree algorithm (in pseudo C code) takes arguments $(D[],n,x,N)$; data array $D[0..n-1]\in[0,1)^n$, a point $x\in\mathbb{R}$, and an integer N. It returns $(p,h,\bar{h},\tilde{p}[],r)$; the logarithmic data evidence $p\hat{=}\ln p(D)$, the expected tree height $h\hat{=}E[h_{\vec{q}_*}(x)|D]$ at x, the average expected tree height $\bar{h}\hat{=}E[\bar{h}_{\vec{q}_*}|D]$, the model dimension distribution $\tilde{p}[0..N-1]\hat{=}P[N_{\vec{q}_*}=..|D]$, and the number of recursive function calls r i.e. the size of the generated tree. s, u and α are the global model parameters. Computation time is about $N^2 n\log n$ nano-seconds on a 1GHz P4 laptop.

BayesTree($D[],n,x,N$)

```
⌈ if (n≤1 and (n==0 or D[0]==x or x∉[0,1)))
  ⌈ if (x∈[0,1)) h=s/u; else h=0;
    h̄=s/u; p=ln(1); r=1;
  ⌊ for(k=0,..,N−1) p̃[k]=a_k;                          /* see (33) */
  else
  ⌈ n_0=n_1=0;
    for(i=0,..,n−1)
    ⌈ if (D[i]<½) then[ D_0[n_0]=2D[i];     n_0=n_0+1;]
    ⌊                 else [D_1[n_1]=2D[i]−1; n_1=n_1+1;]
    (p_0,h_0,h̄_0,p̃_0[],r_0)=BayesTree(D_0[],n_0,2x,N−1);
    (p_1,h_1,h̄_1,p̃_1[],r_1)=BayesTree(D_1[],n_1,2x−1,N−1);
    t=p_0+p_1−lnw(n_0,n_1);                             /* see (14) */
    if (t<100) then p=ln(u+s·exp(t));
               else  p=t+ln(s);
    g=1−u·exp(−p);
    if (x∈[0,1)) then h=g·(1+h_0+h_1); else h=0;
    h̄=g·(1+(n_0+α)/(n+2α) h̄_0+(n_1+α)/(n+2α) h̄_1);
    p̃[0]=1−g;
    for(k=0,..,N−1) p̃[k+1]=g·Σ_{i=0}^{k} p̃_0[i]·p̃_1[k−i];
  ⌊ r=1+r_0+r_1;
⌊ return (p,h,h̄,p̃[],r);
```

How algorithm BayesTree() works. Since evidence $p(D)$ and weight $1/w_n$ can grow exponentially with n, we have to store and use their logarithms. So the algorithm returns $p\hat{=}\ln p(D)$. In the $n\leq 1$ branch, the closed form solutions

$p\widehat{=}\ln p(\emptyset)=\ln(1)$, $h\widehat{=}E[h_{\vec{q}_*}(x)|\emptyset \text{ or } x]=1$, $\bar{h}\widehat{=}E[\bar{h}_{\vec{q}_*}|D]=1$, and $\tilde{p}[k]=a_k$ have been used to truncate the recursion. If $D=(x^1)\neq x$, we have to recurse further until x falls in an empty interval. In this case or if $n>1$ we partition D into points left and right of $\frac{1}{2}$. Then we rescale the points to [0,1) and store them in D_0 and D_1, respectively. Array D could have been reused (like in quick sort) without allocating two new arrays. Then, algorithm BayesTree() is recursively called for each partition. The results are combined according to the recursions derived in Section 2. $\ln w$ can be computed from (14) via $\ln n!=\sum_{k=1}^{n}\ln k$. (Practically, pre-tabulating a_k or $n!$ does not improve overall performance). For computing p we need to use $\ln(\frac{1}{2}(1+e^t))\doteq t-\ln 2$ to machine precision for large t in order to avoid numerical overflow.

Remarks. Strictly speaking, the algorithm has runtime $O(n\log n)$, since the sorting effectively runs once through all data at each level. If we assume that the data are presorted or the counts n_z are given, then the algorithm is $O(n)$ [Hut07].

We have not presented the part handling multi-points. Given the formulas in Section 6, this is easy. The complete C code is available from [Hut07].

Note that x passed to BayesTree() is *not* and cannot be used to compute $p(x|D)$. For this, one has to call BayesTree() twice, with D and (D,x), respectively. The quadratic order in N is due to the convolution, which could be reduced to $O(N\log N)$ by transforming it to a scalar product in Fourier space with FFT.

Multiply calling BayesTree(), e.g. for computing the predictive density function $p(x|D)$ on a fine x-grid, is inefficient. But it is easy to see that if we once pre-compute the evidence $p_z(D_z)$ for all z up to the separation level in time $O(n)$, we can compute "local" quantities like $p(x|D)$ at x in time $O(\log n)$. This is because only the branch containing x needs to be recursed, the other branch is immediately available, since it involves the already pre-computed evidence only. The predictive density $p(x|D)=E[q(x)|D]$ and higher moments, the distribution function $P[x\leq a|D]$, updating D by adding or removing one data item, and most other local quantities can be computed in time $O(\log n)$ by such a linear recursion.

A good way of checking correctness of the implementation *and* of the derived formulas, is to force some *minimal* recursion depth m'. The results must be independent of m', since the closed-form speedups are exact and applicable anywhere beyond the separation level.

8 Numerical Examples

Graphs and examples. To get further insight into the behavior of our model, we numerically investigated some example distributions $\dot{q}()$. We have chosen elementary functions, which can be regarded as prototypes for more realistic functions. They include the Beta, linear, a singular, and piecewise constant distributions with finite and infinite Bayes trees. These examples on [0,1) also shed light on the other spaces discussed in Section 2, since they are isomorphic. The posteriors, model dimensions,

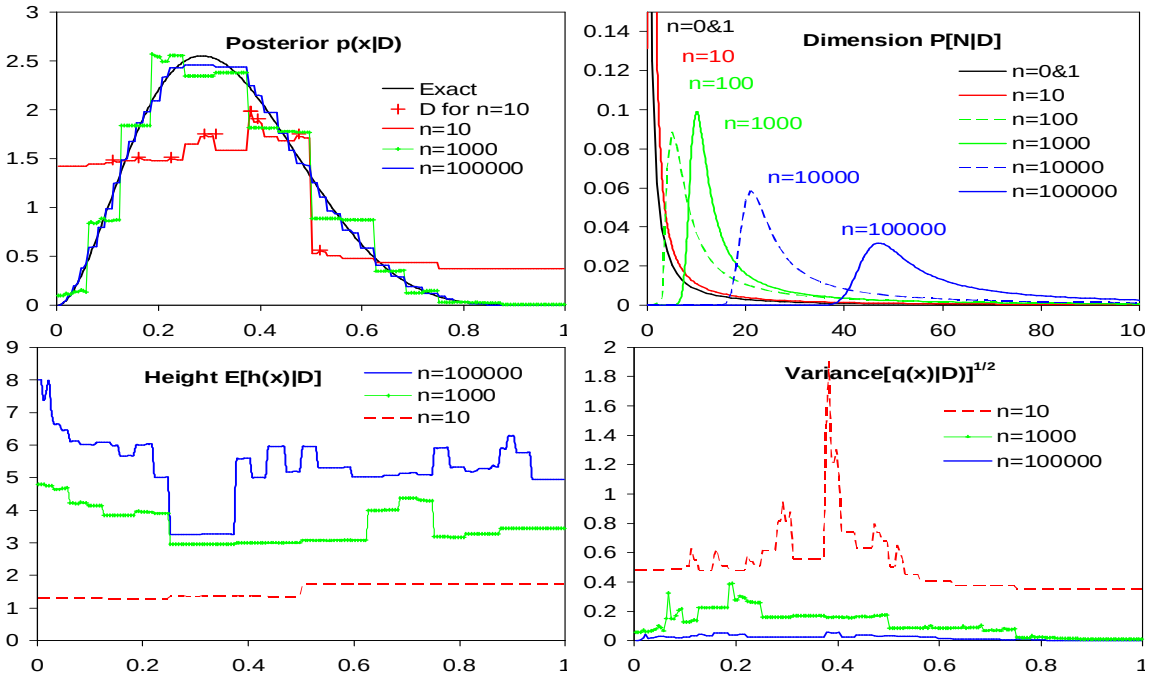

Figure 5: BayesTree() results for the **Beta(3,6)** $\propto x^2(1-x)^5$ distribution, prototype for a smooth distribution.

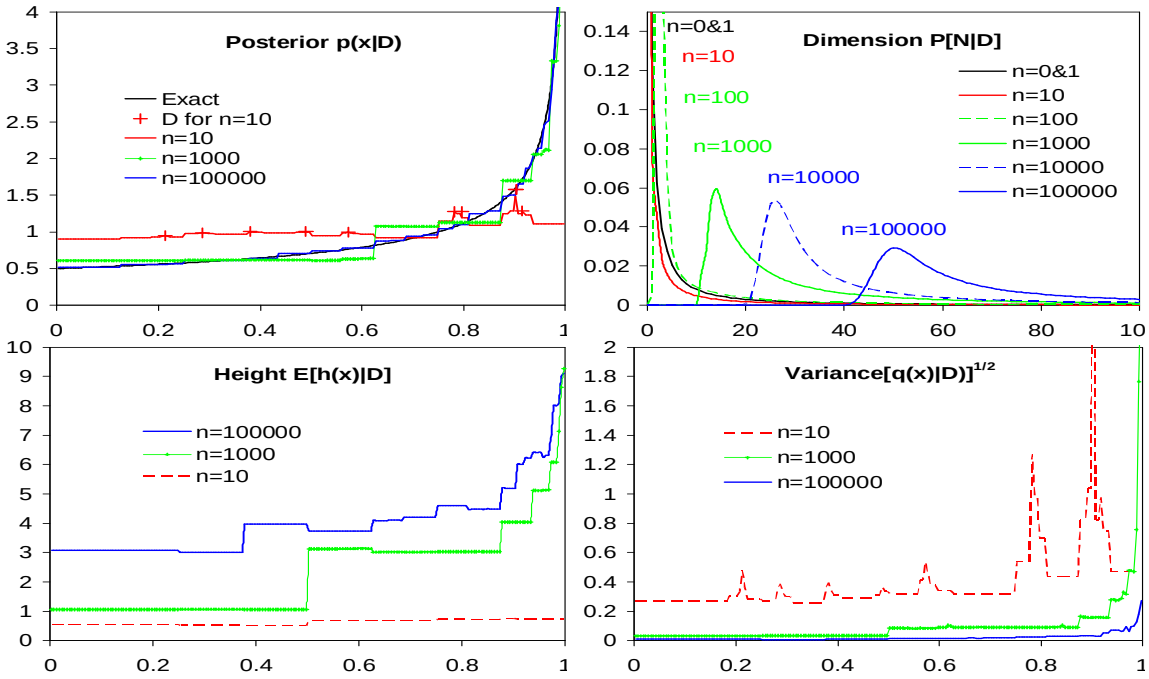

Figure 6: BayesTree() results for the **Singular** distribution $\dot{q}(x)=2/\sqrt{1-x}$, prototype for a proper singular distribution.

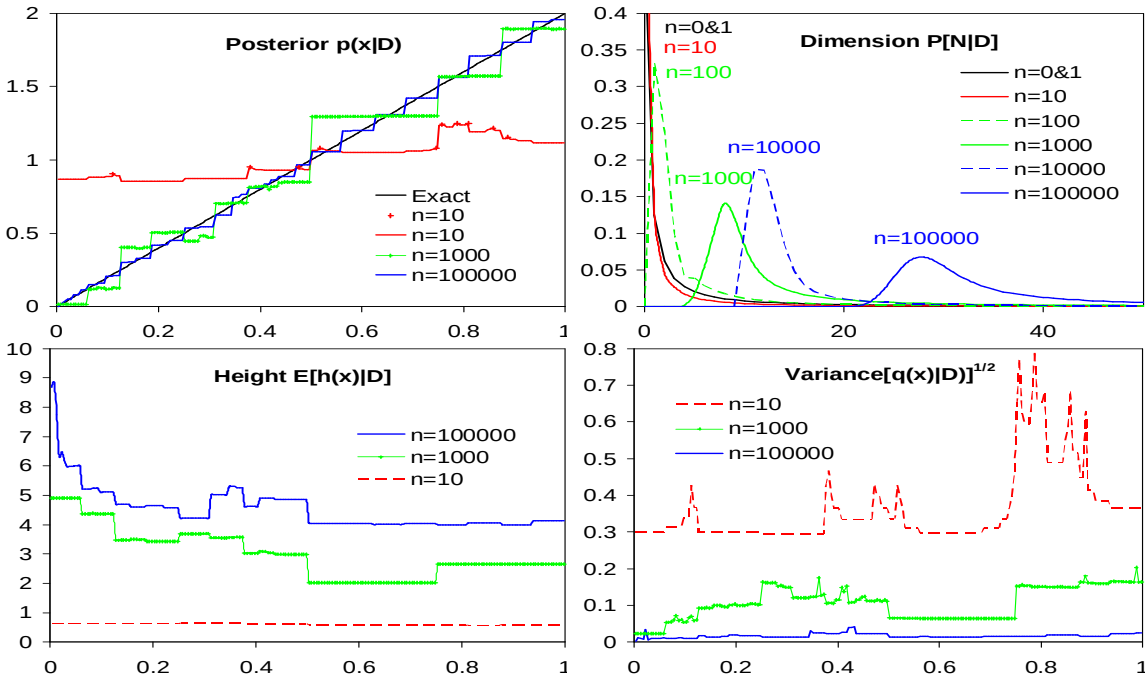

Figure 7: BayesTree() results for the **Linear** distribution $\dot{q}(x) = 2x$, prototype for a continuous function.

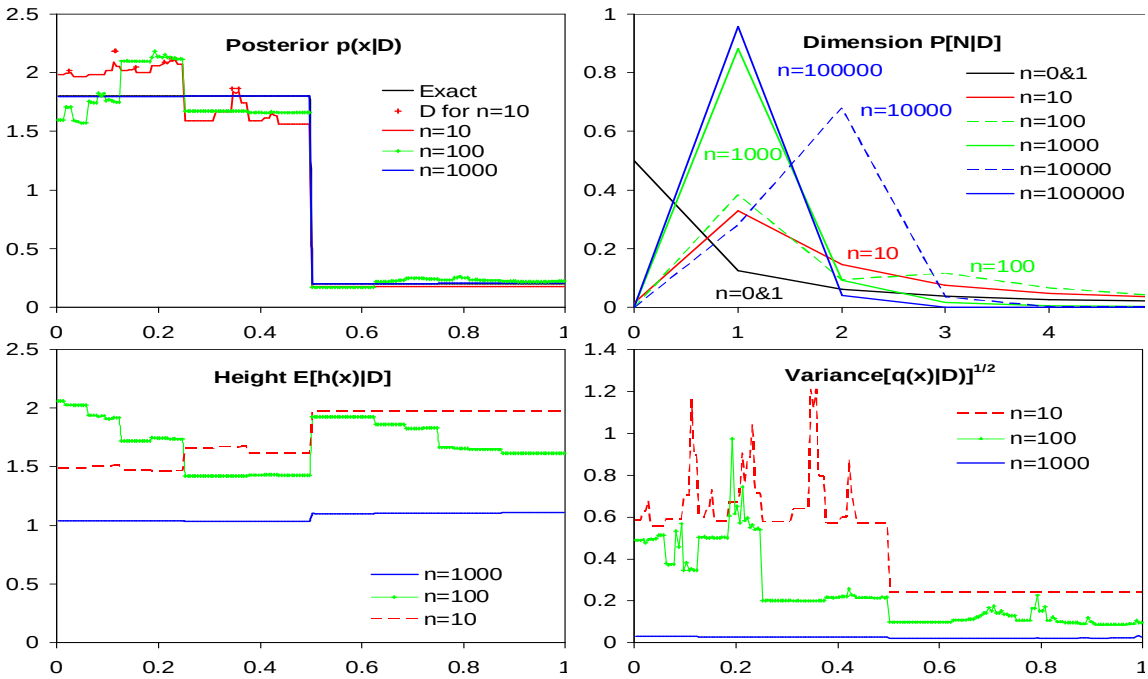

Figure 8: BayesTree() results for the **Jump-at-1/2** distribution $\dot{q}(x) = 9/5$ for $x < 1/2$ and $q(x) = 1/5$ for $x \geq 1/2$, prototype for a piecewise constant function with a finite Bayes tree.

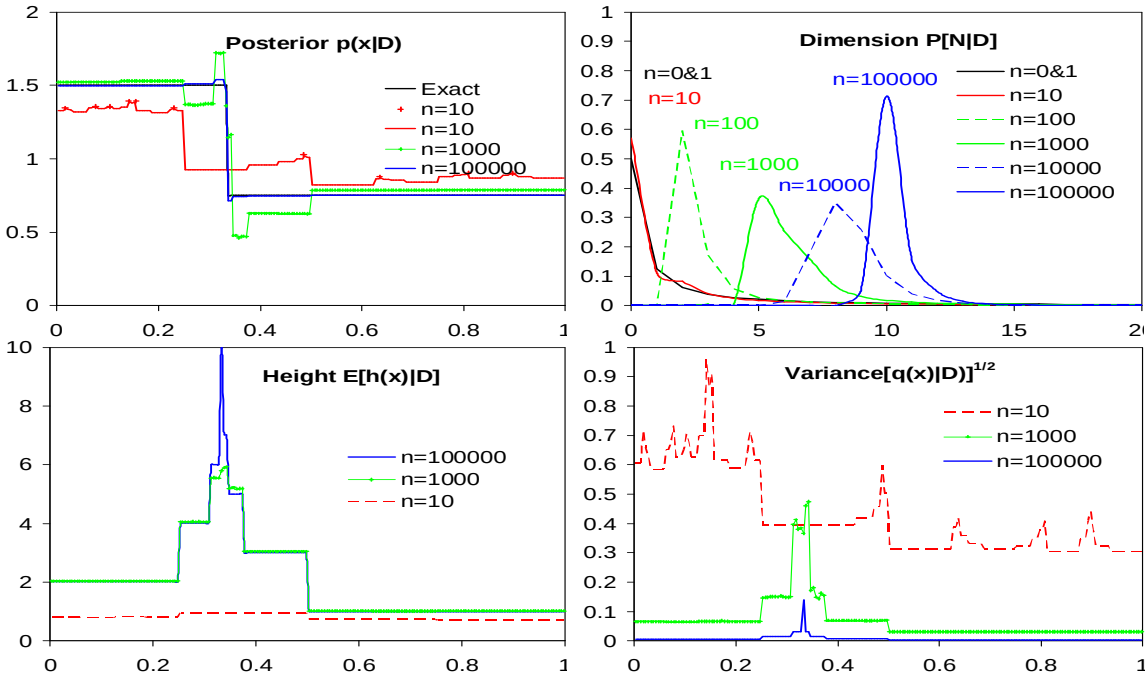

Figure 9: BayesTree() results for the **Jump-at-1/3** distribution $\dot{q}(x) = 2/3$ for $x < \mathbf{1/3}$ and $q(x) = 1/3$ for $x \geq 1/3$, prototype for a piecewise constant function with an infinite Bayes tree.

tree heights, and variances are plotted in Figures 5–9 for random samples D of sizes $n = 10^0,...,10^5$. We first discuss observations common to all sampling distributions, thereafter specific aspects. All experiments were performed with split probability $s = \frac{1}{2}$ and uniform distribution $\alpha = 1$.

General observations. The posteriors $p(x|D)$ clearly converge for $n \to \infty$ to the true distribution $\dot{q}()$, accompanied by a (necessary) moderate growth of the effective dimension (except for Jump-at-1/2). For $n=10$ we show the data points. It is visible how each data point pulls the posterior up, as it should be ("one sample seldom comes alone").

Compare this to an empirical bin model with N bins. Since each bin contains $O(n/N)$ data points, the frequency estimate n_{bin}/n of $\dot{q}(bin)$ has accuracy $O(\sqrt{N/n})$. The minimal error when approximating a continuous function by a piecewise constant function with bin size $1/N$ is $O(1/N)$, so the estimate has total error $\max\{O(\sqrt{N/n}), O(1/N)\}$ with minimum $O(n^{-1/3})$ at[4] $N = n^{1/3}$. This is nicely consistent with our model dimension. Look at the maxima of dimension distribution $P[N|D]$ or count the number of significant jumps in the posterior $p(x|D)$.

The figures also show that the posterior variances $\mathrm{Var}[q(x)|D]$ converge to zero for $n \to \infty$, but diverge when x tends to a point in D, consistent with the theoretical

[4]Sometimes heuristic $N = \sqrt{n}$ is proposed, which makes no sense.

analysis of multi-points in Section 6. The expected tree height $E[h(x)|D]$ at x correctly reflects the local needs for (non)splits.

Specific observations. *Beta:* The Beta distribution $\text{Beta}(x|\alpha,\beta) \propto x^{\alpha-1}(1-x)^{\beta-1}$ is prototypical for a smooth unimodal distribution. Apart from local jitter, the simplest model consistent with data size $n=10$ is essentially a Jump-at-1/2 function (see below). The tree height slowly increases with n with a dip around the "flat" maximum of the Beta, since a constant approximation works well there.

Singular: We used the distribution $\dot{q}(x) = 2/\sqrt{1-x}$ as a prototype for a proper singular distribution. The tree height is necessarily larger near the singularity at $x=1$.

Linear: Once continuously differentiable functions are locally linear, so the linear distribution $\dot{q}(x) = 2x$ serves as a prototype for them. The better approximation of $p(x|D)$ near 0 versus near 1, accompanied by a higher tree, is remarkable. First, there are fewer data points near 0 to warrant this, and second, the region is less interesting, since more samples are at 1. So we expected quite the opposite behavior. We currently have no explanation for this phenomenon.

Jump-at-1/2: Also illustrative are distributions with finite Bayes tree, i.e. piecewise constant functions with discontinuities only at binary fractions. We consider the prototype that jumps at $x=1/2$. All quantities converge rapidly. We see that model dimension and tree height stay finite in this case, as they should. Both converge to the minimal consistent value 1. The variance in the left and right half of [0,1) is roughly proportional to $\dot{q}$ therein.

Jump-at-1/3: A jump at a non-binary fraction cannot be modeled with a finite tree. Convergence is slower than for Jump-at-1/2, but faster than for the other examples, which makes sense since only one branch of the tree has to grow to infinity. This shows up in a slower increase of dimension, a converging height function with singularity at 1/3, and a narrowing spike in the variance.

9 Discussion

We presented a Bayesian model on infinite trees, where we split a node into two subtrees with some probability, and assigned a Beta distributed probability to each subtree. We were primarily interested in the case of zero prior knowledge. In this case, scale invariance and symmetry should be preserved. Scale invariance requires the parameters to be the same in each node and symmetry requires a symmetric Beta, leaving one splitting probability s and one Beta parameter β as adjustable parameters. We devised closed form expressions for various inferential quantities of interest at the data separation level, which led to an exact algorithm with runtime essentially linear in the data size.

We introduced and studied this two-parameter tree-model class. The most interesting case of splitting probability $s=\frac{1}{2}$ and uniform prior over subtrees $\beta=1$ has been studied in more detail. The theoretical and numerical model behavior was

very reasonable, e.g. consistency (no underfitting) and low finite effective dimension (no overfitting). Higher moments can be made finite by smaller s or larger β.

There are various natural generalizations of our model. The splitting probability s and Beta parameter β could be made dependent on the node of the tree, which allows incorporating prior knowledge. k-ary trees could be allowed with Beta generalized to Dirichlet distributions. Non-symmetric partitions are straightforward to implement by replacing all $\delta(q_z-\frac{1}{2})$ with $\delta(q_z-|\Gamma_z|/|\Gamma_{z_{1:l-1}}|)$, and possibly using non-symmetric Betas. The expected entropy can also be computed by allowing fractional counts n_z and noting that $x\ln x=\frac{d}{dx}x^\alpha|_{\alpha=1}$ [Hut02, HZ05]. A sort of maximum a posteriori (MAP) tree skeleton can also easily be read off from (13). A node Γ_z in the MAP-like tree is a leaf iff $s\frac{p_{z0}(D_{z0})p_{z1}(D_{z1})}{w(n_{z0},n_{z1})}<u$. A challenge is to generalize the model from piecewise constant to piecewise linear continuous functions, at least for $\Gamma=[0,1)$. Independence of subtrees no longer holds, which was key in our analysis.

If Γ is not already a tree or binary string, but an interval, a major problem of Polya trees and our tree model are partition artifacts in the estimated density. Numerically but unlikely analytically it is possible to average over boundary locations like in [PRLW03] and smooth out discontinuities. Interestingly for flat bin estimation, analytical averaging *is* possible via dynamic programming [EF05].

References

[Bis06] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[BM98] A. A. Borovkov and A. Moullagaliev. *Mathematical Statistics*. Gordon & Breach, 1998.

[DLR77] A. P. Dempster, N. Laird, and D. Rubin. Maximum likelihood estimation for incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, Series B 39:1–38, 1977.

[EF05] D. Endres and P. Földiák. Bayesian bin distribution inference and mutual information. *IEEE Transactions on Information Theory*, 51(11):3766–3779, 2005.

[EW95] M. Escobar and M. West. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90:577–588, 1995.

[Fab64] J. Fabius. Asymptotic behavior of Bayes estimates. *Annals of Mathematical Statistics*, 35:846–856, 1964.

[Fer73] T. S. Ferguson. On the mathematical foundations of theoretical statistics. *Annals of Statistics*, 1(2):209–230, 1973.

[GM03] A. G. Gray and A. W. Moore. Nonparametric density estimation: Toward computational tractability. In *SIAM International Conf. on Data Mining*, volume 3, 2003.

[Goo83] I. J. Good. Explicativity, corroboration, and the relative odds of hypotheses. In *Good thinking: The Foundations of Probability and its applications*. University of Minnesota Press, Minneapolis, MN, 1983.

[Goo84] I. J. Good. The best explicatum for weight of evidence. *Journal of Statistical Computation and Simulation*, 19:294–299, 1984.

[Hut02] M. Hutter. Distribution of mutual information. In *Advances in Neural Information Processing Systems 14 (NIPS'01)*, pages 399–406, Cambridge, MA, 2002. MIT Press.

[Hut05a] M. Hutter. Fast non-parametric Bayesian inference on infinite trees. In *Proc. 10th International Conf. on Artificial Intelligence and Statistics (AISTATS'05)*, pages 144–151. Society for Artificial Intelligence and Statistics, 2005.

[Hut05b] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2005. 300 pages, http://www.hutter1.net/ai/uaibook.htm.

[Hut07] M. Hutter. Additional material to article, 2007. http://www.hutter1.net/official/bib.htm#bayestreex.

[HZ05] M. Hutter and M. Zaffalon. Distribution of mutual information from complete and incomplete data. *Computational Statistics & Data Analysis*, 48(3):633–657, 2005.

[Jay03] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, MA, 2003.

[Jef35] H. Jeffreys. Some tests of significance, treated by the theory of probability. In *Proc. Cambridge Philosophical Society*, volume 31, pages 203–222, 1935.

[KF98] D. Koller and R. Fratkina. Using learning for approximation in stochastic processes. In *Proc. 15th International Conference on Machine Learning (ICML'98)*, pages 287–295, 1998.

[KK97] A. V. Kozlov and D. Koller. Nonuniform dynamic discretization in hybrid networks. In *Proc. 13th Conf. on Uncertainty in Artificial Intelligence (UAI'97)*, pages 314–325, 1997.

[KM07] P. Kontkanen and P. Myllymki. MDL histogram density estimation. In *Proc. 11th International Conf. on Artificial Intelligence and Statistics (AISTATS'07)*. Society for Artificial Intelligence and Statistics, 2007.

[Lav92] M. Lavine. Some aspects of Polya tree distributions for statistical modelling. *Annals of Statistics*, 20:1222–1235, 1992.

[Lav94] M. Lavine. More aspects of Polya tree distributions for statistical modelling. *Annals of Statistics*, 22:1161–1176, 1994.

[Lem03] J. C. Lemm. *Bayesian Field Theory and Approximate Symmetries*. Johns Hopkins University Press, 2003.

[LLW07] H. Liu, J. Lafferty, and L. Wasserman. Sparse nonparametric density estimation in high dimensions using the rodeo. In *Proc. 11th International Conf. on Artificial Intelligence and Statistics (AISTATS'07)*. Society for Artificial Intelligence and Statistics, 2007.

[Mac03] D. J. C. MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, Cambridge, MA, 2003.

[PRLW03] S. M. Paddock, F. Ruggeri, M. Lavine, and M. West. Randomised Polya tree models for nonparametric Bayesian inference. *Statistica Sinica*, 13(2):443–460, 2003.

[PW02] S. Petrone and L. Wasserman. Consistency of Bernstein polynomial posteriors. *Journal of the Royal Statistical Society*, B 64:79–100, 2002.