

Mondrian Forests: Efficient Online Random Forests

Supplementary Material

A Posterior inference and prediction using the HNSP

Recall that we use a hierarchical Bayesian approach to specify a smooth label distribution $p_T(y|\mathbf{x}, \mathcal{D}_{1:N})$ for each tree T . The label prediction at a test point \mathbf{x} will depend on where \mathbf{x} falls relative to the existing data in the tree T . In this section, we assume that \mathbf{x} lies within one of the leaf nodes in T , i.e., $\mathbf{x} \in B_{\text{leaf}(\mathbf{x})}^x$, where $\text{leaf}(\mathbf{x}) \in \text{leaves}(T)$. If \mathbf{x} does not lie within any of the leaf nodes in T , i.e., $\mathbf{x} \notin \cup_{j \in \text{leaves}(T)} B_j^x$, one could extend the tree by sampling T' from $\text{MTx}(\lambda, T, \mathbf{x})$, such that \mathbf{x} lies within a leaf node in T' and apply the procedure described below using the extended tree T' . Appendix B describes this case in more detail.

Given training data $\mathcal{D}_{1:N}$, a Mondrian tree T and the hierarchical prior over \mathcal{G} , the predictive label distribution $p_T(y|\mathbf{x}, \mathcal{D}_{1:N})$ is obtained by integrating over \mathcal{G} , i.e.

$$p_T(y|\mathbf{x}, \mathcal{D}_{1:N}) = \mathbb{E}_{\mathcal{G} \sim p_T(\mathcal{G}|\mathcal{D}_{1:N})} [G_{\text{leaf}(\mathbf{x}), y}] = \bar{G}_{\text{leaf}(\mathbf{x}), y}.$$

Hence, the prediction is given by $\bar{G}_{\text{leaf}(\mathbf{x})}$, the posterior mean at $\text{leaf}(\mathbf{x})$. The posterior mean $\bar{G}_{\text{leaf}(\mathbf{x})}$ can be computed using existing techniques, which we review in the rest of this section.

Posterior inference in the HNSP is a special case of posterior inference in hierarchical PYP (HPYP). Teh [22] considers the HPYP with multinomial likelihood (in the context of language modeling)—the model considered here (HNSP with multinomial likelihood) is a special case of [22]. Hence, we just sketch the high level picture and refer the reader to [22] for further details. We first describe posterior inference given N data points $\mathcal{D}_{1:N}$ (batch setting), and later explain how to adapt inference to the online setting. Finally, we describe the computation of the predictive posterior distribution.

Batch setting

Posterior inference is done using the Chinese restaurant process representation, wherein every node of the decision tree is a restaurant; the training data points are the customers seated in the tables associated with the leaf node restaurants; these tables are in turn customers at the tables in their corresponding parent level restaurant; the dish served at each table is the class label. Exact inference is intractable and hence we resort to approximations. In particular, we use the approximation known as the interpolated Kneser-Ney (IKN) smoothing, a popular smoothing technique for language modeling [13]. The IKN smoothing can be interpreted as an approximate inference scheme for the HPYP, where the number of tables serving a particular dish in a restaurant is at most one [22]. More precisely, if $c_{j,k}$ denotes the number of customers at restaurant j eating dish k and $\text{tab}_{j,k}$ denotes the number of tables at restaurant j serving dish k , the IKN approximation sets $\text{tab}_{j,k} = \min(c_{j,k}, 1)$. The counts $c_{j,k}$ and $\text{tab}_{j,k}$ can be computed in a single bottom-up pass as follows: for every leaf node $j \in \text{leaves}(T)$, $c_{j,k}$ is simply the number of training data points with label k at node j ; for every internal node $j \in T \setminus \text{leaves}(T)$, we set $c_{j,k} = \text{tab}_{\text{left}(j),k} + \text{tab}_{\text{right}(j),k}$. For a leaf node j , this procedure is summarized in Algorithm 5. (Note that this pseudocode just serves as a reference; in practice, these counts are updated in an online fashion, as described in Algorithm 6.)

Algorithm 5 InitializePosteriorCounts(j)

- 1: For all k , set $c_{jk} = \#\{n \in N(j) : y_n = k\}$
- 2: Initialize $j' = j$
- 3: **while** True **do**
- 4: **if** $j' \notin \text{leaves}(T)$ **then**
- 5: For all k , set $c_{j'k} = \text{tab}_{\text{left}(j'),k} + \text{tab}_{\text{right}(j'),k}$
- 6: For all k , set $\text{tab}_{j'k} = \min(c_{j'k}, 1)$ \triangleright IKN approximation
- 7: **if** $j' = \epsilon$ **then**
- 8: **return**
- 9: **else**
- 10: $j' \leftarrow \text{parent}(j')$

Posterior inference: online setting

It is straightforward to extend inference to the online setting. Adding a new data point $\mathcal{D} = (\mathbf{x}, y)$ affects only the counts along the path from the root to the leaf node of that data point. We update the counts in a bottom-up fashion, starting at the leaf node containing the data point, $\text{leaf}(\mathbf{x})$. Due to the nature of the IKN approximation, we can stop at the internal node j where $c_{j,y} = 1$ and need not traverse up till the root. This procedure is summarized in Algorithm 6.

Algorithm 6 UpdatePosteriorCounts(j, y)

```

1:  $c_{jy} \leftarrow c_{jy} + 1$ 
2: Initialize  $j' = j$ 
3: while True do
4:   if  $\text{tab}_{j'y} = 1$  then                                 $\triangleright$  none of the counts above need to be updated
5:     return
6:   else
7:     if  $j' \notin \text{leaves}(\mathbb{T})$  then
8:        $c_{j'y} = \text{tab}_{\text{left}(j'),y} + \text{tab}_{\text{right}(j'),y}$ 
9:        $\text{tab}_{j'y} = \min(c_{j'y}, 1)$                                  $\triangleright$  IKN approximation
10:      if  $j' = \epsilon$  then
11:        return
12:      else
13:         $j' \leftarrow \text{parent}(j')$ 

```

Predictive posterior computation Given the counts $c_{j,k}$ and table assignments $\text{tab}_{j,k}$, the predictive probability (i.e., posterior mean) at node j can be computed recursively as follows:

$$\bar{G}_{jk} = \begin{cases} \frac{c_{j,k} - d_j \text{tab}_{j,k}}{c_{j,\cdot}} + \frac{d_j \text{tab}_{j,\cdot}}{c_{j,\cdot}} \bar{G}_{\text{parent}(j),k} & c_{j,\cdot} > 0, \\ \bar{G}_{\text{parent}(j),k} & c_{j,\cdot} = 0, \end{cases} \quad (4)$$

where $c_{j,\cdot} = \sum_k c_{j,k}$, $\text{tab}_{j,\cdot} = \sum_k \text{tab}_{j,k}$, and $d_j := \exp(-\gamma(\tau_j - \tau_{\text{parent}(j)}))$ is the *discount* for node j , defined in Section 4. Informally, the discount interpolates between the counts c and the prior. If the discount $d_j \approx 1$, then \bar{G}_j is more like its parent $\bar{G}_{\text{parent}(j)}$. If $d_j \approx 0$, then \bar{G}_j weights the counts more. These predictive probabilities can be computed in a single top-down pass as shown in Algorithm 7.

Algorithm 7 ComputePosteriorPredictiveDistribution(T, \mathcal{G})

```

1:  $\triangleright$  Description of top-down pass to compute posterior predictive distribution given by (4)
2:  $\triangleright$   $\bar{G}_{jk}$  denotes the posterior probability of  $y = k$  at node  $j$ 
3: Initialize the ordered set  $J = \{\epsilon\}$ 
4: while  $J$  not empty do
5:   Pop the first element of  $J$ 
6:   if  $j = \epsilon$  then
7:      $\bar{G}_{\text{parent}(\epsilon)} = H$ 
8:   Set  $d = \exp(-\gamma(\tau_j - \tau_{\text{parent}(j)}))$ 
9:   For all  $k$ , set  $\bar{G}_{jk} = c_{j,\cdot}^{-1} (c_{j,k} - d \text{tab}_{j,k} + d \text{tab}_{j,\cdot} \bar{G}_{\text{parent}(j),k})$ 
10:  if  $j \notin \text{leaves}(\mathbb{T})$  then
11:    Append  $\text{left}(j)$  and  $\text{right}(j)$  to the end of the ordered set  $J$ 

```

B Prediction using Mondrian tree

Let \mathbf{x} denote a test data point. We are interested in the predictive probability of y at \mathbf{x} , denoted by $p_T(y|\mathbf{x}, \mathcal{D}_{1:N})$. As in typical decision trees, the process involves a top-down tree traversal, starting from the root. If \mathbf{x} is already ‘contained’ in the tree T , i.e., if $\mathbf{x} \in B_j^x$ for some leaf $j \in \text{leaves}(\mathbb{T})$, then the prediction is taken to be $\bar{G}_{\text{leaf}(\mathbf{x})}$, which is computed as described in Appendix A. Otherwise,

we somehow need to incorporate \mathbf{x} . One choice is to extend T by sampling T' from $\text{MTx}(\lambda, T, \mathbf{x})$ as described in Algorithm 3, and set the prediction to \bar{G}_j , where $j \in \text{leaves}(T')$ is the leaf node containing \mathbf{x} . A particular extension T' might lead to an overly confident prediction; hence, we average over *every* possible extension T' . This expectation can be carried out analytically, using properties of the Mondrian process, as we show below.

Let $\text{ancestors}(j)$ denote the set of all ancestors of node j . Let $\text{path}(j) = \{j\} \cup \text{ancestors}(j)$, that is, the set of all nodes along the ancestral path from j to the root. Recall that $\text{leaf}(\mathbf{x})$ is the unique leaf node in T such that $\mathbf{x} \in B_{\text{leaf}(\mathbf{x})}$. If the test point $\mathbf{x} \in B_{\text{leaf}(\mathbf{x})}^x$ (i.e., \mathbf{x} lies within the ‘gray rectangle’ at the leaf node), it can never branch off; else, it can branch off at one or more points along the path from the root to $\text{leaf}(\mathbf{x})$. More precisely, if \mathbf{x} lies outside B_j^x at node j , the probability that \mathbf{x} will branch off into its own node at node j , denoted by⁶ $p_j^s(\mathbf{x})$, is equal to the probability that a split exists in B_j outside B_j^x , which is

$$p_j^s(\mathbf{x}) = 1 - \exp(-\Delta_j \eta_j(\mathbf{x})), \quad \text{where } \eta_j(\mathbf{x}) = \sum_d (\max(x_d - u_{jd}^x, 0) + \max(\ell_{jd}^x - x_d, 0)),$$

and $\Delta_j = \tau_j - \tau_{\text{parent}(j)}$. Note that $p_j^s(\mathbf{x}) = 0$ if \mathbf{x} lies within B_j^x (i.e., if $\ell_{jd}^x \leq x_d \leq u_{jd}^x$ for all d). The probability of \mathbf{x} not branching off before reaching node j is given by $\prod_{j' \in \text{ancestors}(j)} (1 - p_{j'}^s(\mathbf{x}))$.

If $\mathbf{x} \in B_{\text{leaf}(\mathbf{x})}^x$, the prediction is given by $\bar{G}_{\text{leaf}(\mathbf{x})}$. If there is a split in B_j outside B_j^x , let \tilde{j} denote the new parent of j and $\text{child}(\tilde{j})$ denote the child node containing just the test data point; in this case, the prediction is $\bar{G}_{\text{child}(\tilde{j})}$. Averaging over the location where the test point branches off, we obtain

$$p_T(y|\mathbf{x}, \mathcal{D}_{1:N}) = \sum_{j \in \text{path}(\text{leaf}(\mathbf{x}))} \left(\prod_{j' \in \text{ancestors}(j)} (1 - p_{j'}^s(\mathbf{x})) \right) F_j(\mathbf{x}), \quad (5)$$

where

$$F_j(\mathbf{x}) = p_j^s(\mathbf{x}) \mathbb{E}_{\Delta_j} [\bar{G}_{\text{child}(\tilde{j})}] + \mathbb{1}[j = \text{leaf}(\mathbf{x})] (1 - p_j^s(\mathbf{x})) \bar{G}_{\text{leaf}(\mathbf{x})}. \quad (6)$$

The second term in $F_j(\mathbf{x})$ needs to be computed only for the leaf node $\text{leaf}(\mathbf{x})$ and is simply the posterior mean of $G_{\text{leaf}(\mathbf{x})}$ weighted by $1 - p_{\text{leaf}(\mathbf{x})}^s(\mathbf{x})$. The posterior mean of $G_{\text{leaf}(\mathbf{x})}$, given by $\bar{G}_{\text{leaf}(\mathbf{x})}$, can be computed using (4). The first term in $F_j(\mathbf{x})$ is simply the posterior mean of $G_{\text{child}(\tilde{j})}$, averaged over $\Delta_{\tilde{j}}$, weighted by $p_{\tilde{j}}^s(\mathbf{x})$. Since no labels are observed in $\text{child}(\tilde{j})$, $c_{\text{child}(\tilde{j})} = 0$, hence from (4), we have $\bar{G}_{\text{child}(\tilde{j})} = \bar{G}_{\tilde{j}}$. We compute $\bar{G}_{\tilde{j}}$ using (4). We average over $\Delta_{\tilde{j}}$ due to the fact that the discount in (4) for the node \tilde{j} depends on $\tau_{\tilde{j}} - \tau_{\text{parent}(\tilde{j})} = \Delta_{\tilde{j}}$. To average over all valid split times $\tau_{\tilde{j}}$, we compute expectation w.r.t. $\Delta_{\tilde{j}}$ which is distributed according to a truncated exponential with rate $\eta_j(\mathbf{x})$, truncated to the interval $[0, \Delta_j]$.

The procedure for computing $p_T(y|\mathbf{x}, \mathcal{D}_{1:N})$ for any $\mathbf{x} \in \mathbb{R}^D$ is summarized in Algorithm 8. The predictive probability assigned by a Mondrian forest is the average of the predictive probability of the M trees, i.e., $\frac{1}{M} \sum_m p_{T_m}(y|\mathbf{x}, \mathcal{D}_{1:N})$.

C Computational complexity

We discuss the computational complexity associated with a single Mondrian tree. The complexity of a forest is simply M times that of a single tree; however, this computation can be trivially parallelized since there is no interaction between the trees. Assume that the N data points are processed one by one. Assuming the data points form a balanced binary tree after each update, the computational cost of processing the n^{th} data point is at most $\mathcal{O}(\log n)$ (add the data point into its own leaf, update posterior counts for HNSP in bottom-up pass from leaf to root). The overall cost to process N data points is $\mathcal{O}(\sum_{n=1}^N \log n) = \mathcal{O}(\log N!)$, which for large N tends to $\mathcal{O}(N \log N)$ (using Stirling approximation for the factorial function). For offline RF and ERT, the expected complexity with n data points is $\mathcal{O}(n \log n)$. The complexity of the re-trained version is $\mathcal{O}(\sum_{n=1}^N n \log n) = \mathcal{O}(\log \prod_{n=1}^N n^n)$, which for large N tends to $\mathcal{O}(N^2 \log N)$ (using asymptotic expansion of the hyper factorial function).

⁶The superscript s in $p_j^s(\mathbf{x})$ is used to denote the fact that this split ‘separates’ the test data point \mathbf{x} into its own leaf node.

Algorithm 8 Predict(T, \mathbf{x})

- 1: \triangleright Description of prediction using a Mondrian tree, given by (5)
- 2: Initialize $j = \epsilon$ and $p_{\text{NotSeparatedYet}} = 1$
- 3: Initialize $\mathbf{s} = \mathbf{0}_K$ $\triangleright \mathbf{s}$ is K -dimensional vector where $s_k = p_T(y = k | \mathbf{x}, \mathcal{D}_{1:N})$
- 4: **while** True **do**
- 5: Set $\Delta_j = \tau_j - \tau_{\text{parent}(j)}$ and $\eta_j(\mathbf{x}) = \sum_d (\max(x_d - u_{jd}^x, 0) + \max(\ell_{jd}^x - x_d, 0))$
- 6: Set $p_j^s(\mathbf{x}) = 1 - \exp(-\Delta_j \eta_j(\mathbf{x}))$
- 7: **if** $p_j^s(\mathbf{x}) > 0$ **then**
- 8: \triangleright Let \mathbf{x} branch off into its own node $\text{child}(\tilde{j})$, creating a new node \tilde{j} which is the parent of j and $\text{child}(\tilde{j})$. $\overline{G}_{\text{child}(\tilde{j})} = \overline{G}_{\tilde{j}}$ from (4) since $c_{\text{child}(\tilde{j}), \cdot} = 0$.
- 9: Compute expected discount $\bar{d} = \mathbb{E}_{\Delta}[\exp(-\gamma\Delta)]$ where Δ is drawn from a truncated exponential with rate $\eta_j(\mathbf{x})$, truncated to the interval $[0, \Delta_j]$.
- 10: For all k , set $c_{\tilde{j},k} = \text{tab}_{\tilde{j},k} = \min(c_{j,k}, 1)$
- 11: For all k , set $\overline{G}_{\tilde{j}k} = c_{\tilde{j},k}^{-1} (c_{\tilde{j},k} - \bar{d} \text{tab}_{\tilde{j},k} + \bar{d} \text{tab}_{\tilde{j},\cdot} \overline{G}_{\text{parent}(\tilde{j}),k})$ \triangleright Algorithm 7 and (6)
- 12: For all k , update $s_k \leftarrow s_k + p_{\text{NotSeparatedYet}} p_j^s(\mathbf{x}) \overline{G}_{\tilde{j}k}$
- 13: **if** $j \in \text{leaves}(T)$ **then**
- 14: For all k , update $s_k \leftarrow s_k + p_{\text{NotSeparatedYet}} (1 - p_j^s(\mathbf{x})) \overline{G}_{jk}$ \triangleright Algorithm 7 and (6)
- 15: **return** predictive probability \mathbf{s} where $s_k = p_T(y = k | \mathbf{x}, \mathcal{D}_{1:N})$
- 16: **else**
- 17: $p_{\text{NotSeparatedYet}} \leftarrow p_{\text{NotSeparatedYet}} (1 - p_j^s(\mathbf{x}))$
- 18: **if** $x_{\delta_j} \leq \xi_j$ **then** $j \leftarrow \text{left}(j)$ **else** $j \leftarrow \text{right}(j)$ \triangleright recurse to the child where \mathbf{x} lies

D Pseudocode for paused Mondrians

In this section, we discuss versions of SampleMondrianBlock and ExtendMondrianBlock for paused Mondrians. For completeness, we also provide the updates necessary for the IKN approximation.

Algorithm 9 SampleMondrianBlock($j, \mathcal{D}_{N(j)}, \lambda$) version that depends on labels

- 1: Add j to T
- 2: For all d , set $\ell_{jd}^x = \min(\mathbf{X}_{N(j),d})$, $u_{jd}^x = \max(\mathbf{X}_{N(j),d})$ \triangleright dimension-wise min and max
- 3: **if** AllLabelsIdentical($Y_{N(j)}$) **then**
- 4: Set $\tau_j = \lambda$ \triangleright pause Mondrian
- 5: **else**
- 6: Sample E from exponential distribution with rate $\sum_d (u_{jd}^x - \ell_{jd}^x)$
- 7: Set $\tau_j = \tau_{\text{parent}(j)} + E$
- 8: **if** $\tau_j < \lambda$ **then**
- 9: Sample split dimension δ_j with probability of choosing d proportional to $u_{jd}^x - \ell_{jd}^x$
- 10: Sample split location ξ_j along dimension δ_j from a uniform distribution over $\mathcal{U}[\ell_{jd}^x, u_{jd}^x]$
- 11: Set $N(\text{left}(j)) = \{n \in N(j) : \mathbf{X}_{n,\delta_j} \leq \xi_j\}$ and $N(\text{right}(j)) = \{n \in N(j) : \mathbf{X}_{n,\delta_j} > \xi_j\}$
- 12: SampleMondrianBlock($\text{left}(j), \mathcal{D}_{N(\text{left}(j))}, \lambda$)
- 13: SampleMondrianBlock($\text{right}(j), \mathcal{D}_{N(\text{right}(j))}, \lambda$)
- 14: **else**
- 15: Set $\tau_j = \lambda$ and add j to $\text{leaves}(T)$ \triangleright j is a leaf node
- 16: InitializePosteriorCounts(j) \triangleright Algorithm 5

Algorithm 10 ExtendMondrianBlock($T, \lambda, j, \mathcal{D}$) version that depends on labels

```

1: if AllLabelsIdentical( $Y_{N(j)}$ ) then ▷ paused Mondrian leaf
2:   Update extent  $\ell_j^x \leftarrow \min(\ell_j^x, \mathbf{x})$ ,  $\mathbf{u}_j^x \leftarrow \max(\mathbf{u}_j^x, \mathbf{x})$ 
3:   Append  $\mathcal{D}$  to  $\mathcal{D}_{N(j)}$  ▷ append  $\mathbf{x}$  to  $X_{N(j)}$  and  $y$  to  $Y_{N(j)}$ 
4:   if  $y = \text{unique}(Y_{N(j)})$  then
5:     UpdatePosteriorCounts( $j, y$ ) ▷ Algorithm 6
6:     return ▷ continue pausing
7:   else
8:     Remove  $j$  from leaves( $T$ )
9:     SampleMondrianBlock( $j, \mathcal{D}_{N(j)}, \lambda$ ) ▷ un-pause Mondrian
10: else
11:   Set  $\mathbf{e}^\ell = \max(\ell_j^x - \mathbf{x}, 0)$  and  $\mathbf{e}^u = \max(\mathbf{x} - \mathbf{u}_j^x, 0)$  ▷  $\mathbf{e}^\ell = \mathbf{e}^u = \mathbf{0}_D$  if  $\mathbf{x} \in B_j^x$ 
12:   Sample  $E$  from exponential distribution with rate  $\sum_d (e_d^\ell + e_d^u)$ 
13:   if  $\tau_{\text{parent}(j)} + E < \tau_j$  then ▷ introduce new parent for node  $j$ 
14:     Create new Mondrian block  $\tilde{j}$  where  $\ell_{\tilde{j}}^x = \min(\ell_j^x, \mathbf{x})$  and  $\mathbf{u}_{\tilde{j}}^x = \max(\mathbf{u}_j^x, \mathbf{x})$ 
15:     Sample  $\delta_{\tilde{j}}$  with  $\text{Pr}(\delta_{\tilde{j}} = d)$  proportional to  $e_d^\ell + e_d^u$ 
16:     if  $x_{\delta_{\tilde{j}}} > u_{j, \delta_{\tilde{j}}}^x$ , then sample  $\xi_{\tilde{j}}$  from  $\mathcal{U}[u_{j, \delta_{\tilde{j}}}^x, x_{\delta_{\tilde{j}}}]$ , else sample  $\xi_{\tilde{j}}$  from  $\mathcal{U}([x_{\delta_{\tilde{j}}}, \ell_{j, \delta_{\tilde{j}}}^x])$ 
17:     if  $j = \epsilon$  then ▷ set  $\tilde{j}$  as the new root
18:        $\epsilon \leftarrow \tilde{j}$ 
19:     else ▷ set  $\tilde{j}$  as child of parent( $j$ )
20:       if  $j = \text{left}(\text{parent}(j))$ , then  $\text{left}(\text{parent}(j)) \leftarrow \tilde{j}$ , else  $\text{right}(\text{parent}(j)) \leftarrow \tilde{j}$ 
21:       if  $x_{\delta_{\tilde{j}}} > \xi_{\tilde{j}}$  then
22:         Set  $\text{left}(\tilde{j}) = j$  and SampleMondrianBlock( $\text{right}(\tilde{j}), \mathcal{D}, \lambda$ ) ▷ create new leaf for  $x$ 
23:       else
24:         Set  $\text{right}(\tilde{j}) = j$  and SampleMondrianBlock( $\text{left}(\tilde{j}), \mathcal{D}, \lambda$ ) ▷ create new leaf for  $x$ 
25:     else
26:       Update  $\ell_j^x \leftarrow \min(\ell_j^x, \mathbf{x})$ ,  $\mathbf{u}_j^x \leftarrow \max(\mathbf{u}_j^x, \mathbf{x})$  ▷ update extent of node  $j$ 
27:       if  $j \notin \text{leaves}(T)$  then ▷ return if  $j$  is a leaf node, else recurse down the tree
28:         if  $x_{\delta_j} \leq \xi_j$  then  $\text{child}(j) = \text{left}(j)$  else  $\text{child}(j) = \text{right}(j)$ 
29:         ExtendMondrianBlock( $T, \lambda, \text{child}(j), \mathcal{D}$ ) ▷ recurse on child containing  $x$ 

```

E Depth of trees

We computed the average depth of the trees in the forest, where depth of a leaf node is weighted by fraction of data points at that leaf node. The hyper-parameter settings and experimental setup are described in Section 7. Table 1 reports the average depth (and standard deviations) for Mondrian forests trained on different datasets. The values suggest that the depth of the forest scales as $\log N$ rather than N .

Dataset	N_{train}	$\log_2 N_{\text{train}}$	depth
<i>usps</i>	7291	12.8	19.1 ± 1.3
<i>satimages</i>	3104	11.6	17.4 ± 1.6
<i>letter</i>	15000	13.9	23.2 ± 1.8
<i>dna</i>	1400	10.5	12.0 ± 0.3

Table 1: Average depth of Mondrian forests trained on different datasets.

F Comparison to dynamic trees

Dynamic trees [21] approximate the Bayesian posterior over decision trees in an online fashion. Specifically, dynamic trees maintain a particle approximation to the true posterior; the prediction at a test point is a weighted average of the predictions made by the individual particles. While this averaging procedure appears similar to online random forests at first sight, there is a key difference: MF (and other random forests) performs ensemble model combination whereas dynamic trees use Bayesian model averaging. In the limit of infinite data, the Bayesian posterior would converge to a

single tree [15], whereas MF would still average predictions over multiple trees. Hence, we expect MF to outperform dynamic trees in scenarios where a single decision tree is insufficient to explain the data.

To experimentally validate our hypothesis, we evaluate the empirical performance of dynamic trees using the `dynaTree`⁷ R package provided by the authors of the paper. Note that while dynamic trees can use ‘linear leaves’ (strong since prediction at a leaf depends on X) or ‘constant leaves’ for regression tasks, they use ‘multinomial leaves’ for classification tasks which corresponds to a ‘weak learner’. We set the number of particles to 100 (equals the number of trees used in MF) and the number of passes, $R = 2$ (their code does not support $R = 1$) and set the remaining parameters to their default values. Fig. 4 compares the performance of dynamic trees to MF and other random forest variants. (The performance of all methods other than dynamic trees is identical to that of Fig. 3.)

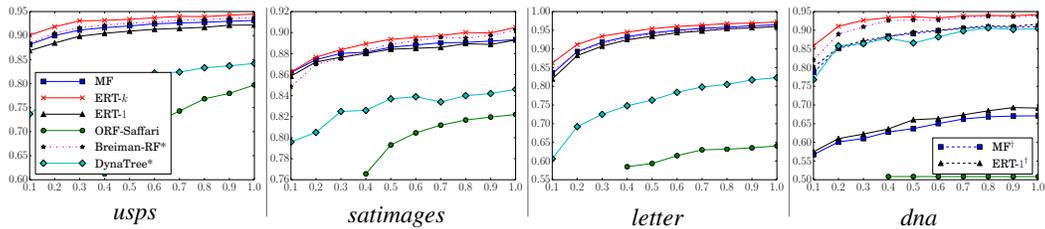


Figure 4: Results on various datasets: y -axis is test accuracy in both rows. x -axis is fraction of training data. The setup is identical to that of Fig. 3. MF achieves significantly higher test accuracies than dynamic trees on *usps*, *satimages* and *letter* datasets and MF[†] achieves similar test accuracy as dynamic trees on the *dna* dataset.

We observe that MF achieves significantly higher test accuracies than dynamic trees on *usps*, *satimages* and *letter* datasets. On *dna* dataset, dynamic trees outperform MF (indicating the usefulness of using labels to guide splits) — however, MF with feature selection (MF[†]) achieves similar performance as dynamic trees. All the batch random forest methods are superior to dynamic trees which suggests that decision trees are not sufficient to explain these real world datasets and that model combination is helpful.

⁷<http://cran.r-project.org/web/packages/dynaTree/index.html>