

bartMachine: A Powerful Tool for Machine Learning

Adam Kapelner* and Justin Bleich*

The Wharton School of the University of Pennsylvania

December 5, 2013

Abstract

We present a new package in R implementing Bayesian Additive Regression Trees (BART). The package introduces many new features for data analysis using BART such as variable selection, interaction detection, model diagnostic plots, incorporation of missing data and the ability to save trees for future prediction. It is significantly faster than the current R implementation, parallelized, and capable of handling both large sample sizes and high-dimensional data.

1 Introduction

Ensemble-of-trees methods have become popular choices for forecasting in both regression and classification problems. Algorithms such as Random Forests (Breiman, 2001) and stochastic gradient boosting (Friedman, 2002) are two well-established and widely employed procedures. Recent advances in ensemble methods include Dynamic Trees (Taddy et al., 2011) and Bayesian Additive Regression Trees (BART, Chipman et al., 2010), which depart from predecessors by being based upon an underlying Bayesian probability model rather than a pure algorithm. BART has demonstrated substantial promise in a wide variety of simulations and real world applications such as predicting avalanches on mountain roads (Blattenberger and Fowles, 2014), predicting how transcription factors interact with DNA (Zhou and Liu, 2008) and predicting movie box office revenues (Eliashberg, 2010). We introduce **bartMachine**, a new R package that significantly expands the capabilities of using BART for data analysis.

Currently, there exists a single implementation of BART on CRAN: **BayesTree**, the package developed by the algorithm's original authors. One of the major drawbacks of this implementation is its lack of a **predict** function. Test data must be provided as an argument during the training phase of the model. Hence, it is impossible to generate forecasts on future data without re-fitting with the entire model. Since the run time is not trivial, forecasting becomes an arduous exercise. A significantly faster implementation of BART that

*Both authors contributed equally to this work.

contains master-slave parallelization exists as Pratola et al. (2013), but this is only available as standalone C++ source code and not integrated with R.

Our goal is to provide a fast, easy-to-use, visualization-rich machine learning package for R users. Our implementation of BART is in Java and is integrated into R via rJava (Urbanek, 2011). From a runtime perspective, our algorithm is significantly faster and is parallelized, allowing computation on as many cores as desired. Not only is the model construction itself parallelized, but the additional features such as prediction, variable selection, and many others can be split across cores as well.

Additionally, we include a variety of expanded and new features. We implement the ability to save trees in memory and provide convenience functions for prediction on test data. We also include plotting functions for both credible and predictive intervals and plots for visually inspecting convergence of BART’s Gibbs sampler. We expand variable importance exploration to include permutation tests and interaction detection. We implement recently developed features for BART including a principled approach to variable selection and the ability to incorporate in prior information for covariates (Bleich et al., 2013). We also implement the strategy found in Kapelner and Bleich (2013) to incorporate missing data during training and handle missingness during prediction.

In Section 2, we provide an overview of BART with a special emphasis on the features that have been extended. In Section 3 we provide a general introduction to the package, highlighting the novel features. Section 4 provides step-by-step examples of the regression capabilities and Section 5 introduces additional step-by-step examples of features unique to classification problems. We conclude in Section 6. Our implementation is slightly different from that in the BayesTree package. We discuss these details more thoroughly in Appendix A and offer predictive performance comparisons in Appendix B.

2 Overview of BART

BART provides a unique approach to nonparametric function estimation using regression trees. Regression trees rely on recursive binary partitioning of predictor space into a set of hyperrectangles in order to approximate some unknown function f . Such trees have received praise for their ability to flexibly fit interactions and nonlinearities. Models composed of sums of regression trees are able to capture additive effects in f better than single trees.

BART can primarily be considered a sum-of-trees ensemble, with a novel estimation approach which relies on a fully Bayesian probability model. Specifically, the BART model can be expressed as:

$$\mathbf{Y} = f(\mathbf{X}) + \boldsymbol{\varepsilon} \approx \mathfrak{T}_1^{\mathcal{L}}(\mathbf{X}) + \mathfrak{T}_2^{\mathcal{L}}(\mathbf{X}) + \dots + \mathfrak{T}_m^{\mathcal{L}}(\mathbf{X}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{I}_n) \quad (1)$$

Here we have m distinct regression trees, each composed of a tree structure, denoted by \mathfrak{T} , and the parameters at the terminal nodes (also called leaves), denoted by \mathcal{L} . The two together, denoted as $\mathfrak{T}^{\mathcal{L}}$ represents an entire tree with both its structure and set of leaf parameters.

The structure of a given tree \mathfrak{T}_t includes information on how any observation recurses down the tree. For each nonterminal (internal) node of the tree, there is a “splitting rule”

taking the form $\mathbf{x}_j < c$ consisting of the “splitting variable” \mathbf{x}_j and the “splitting value” c . An observation moves to the left daughter node if the condition set by the splitting rule is satisfied and to the right daughter node otherwise. The process continues until a terminal node is reached. Then, the observation receives the leaf value of the terminal node as its predicted value. We denote the set of tree’s leaf parameters as $\mathcal{L}_t = \{\mu_{t,1}, \mu_{t,2}, \dots, \mu_{t,b_t}\}$ where b_t is the number of terminal nodes for a given tree.

BART can be distinguished from other ensemble-of-trees models due to its underlying probability model. As a Bayesian model, **BART** consists of a set of priors for the structure and the leaf parameters and a likelihood for data in the terminal nodes. The aim of the priors is to provide regularization, preventing any single regression tree from dominating the total fit.

We provide an overview of the **BART** priors and likelihood and then discuss how draws from the posterior distribution are made. A more complete exposition can be found in Chipman et al. (2010).

2.1 Priors and Likelihood

There are three priors for the **BART** model: a prior on the tree structure itself, a prior on the leaf parameters, and a prior on the error variance σ^2 . The prior on σ^2 is independent from the other two and each tree is independent, yielding:

$$\begin{aligned} \mathbb{P}\left(\mathfrak{F}_1^{\mathcal{L}}, \dots, \mathfrak{F}_m^{\mathcal{L}}, \sigma^2\right) &= \left[\prod_t \mathbb{P}\left(\mathfrak{F}_t^{\mathcal{L}}\right) \right] \mathbb{P}\left(\sigma^2\right) = \left[\prod_t \mathbb{P}\left(\mathcal{L}_t \mid \mathfrak{F}_t\right) \mathbb{P}\left(\mathfrak{F}_t\right) \right] \mathbb{P}\left(\sigma^2\right) \\ &= \left[\prod_t \prod_{\ell} \mathbb{P}\left(\mu_{t,\ell} \mid \mathfrak{F}_t\right) \mathbb{P}\left(\mathfrak{F}_t\right) \right] \mathbb{P}\left(\sigma^2\right) \end{aligned}$$

where the last line follows from an additional assumption of conditional independence of the leaf parameters given the tree’s structure.

The first prior is on the locations of nodes within the tree. Nodes at depth d are non-terminal with probability $\alpha(1+d)^{-\beta}$ where $\alpha \in (0, 1)$ and $\beta \in [0, \infty]$. The purpose of this prior is to keep the trees shallow, thereby limiting the amount of complexity in any single tree. Default values for these hyperparameters of $\alpha = 0.95$ and $\beta = 2$ are recommended by Chipman et al. (2010).

For nonterminal nodes, splitting rules have the following prior. First, a predictor is randomly selected to serve as the splitting variable. In the original formulation, each available predictor is equally likely to be chosen, but this is relaxed in our implementation to allow an arbitrary discrete distribution (see Section 4.11). Then, the splitting value is selected by randomly choosing a value of the selected predictor with equal probability.

The third prior is on the leaf parameters. Given a tree with a set of terminal nodes, each terminal node (or leaf) has a continuous parameter (the leaf parameter) representing the “best guess” of the response in this partition of predictor space. This parameter is the fitted value assigned to any observation that lands in that node. The prior on each of the leaf parameters is given as: $\mu_{\ell} \stackrel{iid}{\sim} \mathcal{N}\left(\mu_{\mu}, \sigma_{\mu}^2\right)$. The expectation, μ_{μ} , is picked to be the range center, $(y_{\min} + y_{\max})/2$. The variance is empirically chosen so that the range center

plus or minus $k = 2$ variances cover 95% of the provided response values in the training set (by default). The aim of this prior is to provide model regularization by shrinking the leaf parameters towards the center of the distribution of the response.

The final prior is on the error variance and is chosen to be $\sigma^2 \sim \text{InvGamma}(\nu/2, \nu\lambda/2)$. λ is determined from the data so that there is a $q = 90\%$ a priori chance (by default) that the **BART** model will improve upon the RMSE from an ordinary least squares regression. Therefore, the majority of the prior probability mass lies below the RMSE from least squares regression. Additionally, the prior is selected so that substantial probability mass is not placed on small values of σ^2 as one aim of this prior is to force more noise into the procedure and prevent overfitting if the model begins to fit too snugly.

Note that the adjustable hyperparameters are α , β , k , ν and q . Default values generally provide good performance, but optimal tuning can be achieved via cross-validation, an automatic feature implemented and described in Section 4.2.

Along with a set of priors, the **BART** a likelihood of responses in the leaves. They are assumed a priori normal with the mean being the “best guess” in the leaf at the current moment (in the Gibbs sampler) and variance being the “best guess” of the variance at the moment i.e. $\mathbf{y}_\ell \sim \mathcal{N}(\mu_\ell, \sigma^2/m)$. Note that σ^2 is scaled by the number of trees in order that the sum of the variances across the m trees is σ^2 .

2.2 Posterior Distribution and Prediction

A Gibbs sampler (Geman and Geman, 1984) is employed to generate draws from the posterior distribution of $\mathbb{P}(\mathfrak{F}_1, \dots, \mathfrak{F}_m, \sigma^2 \mid \mathbf{y})$. A key feature of the Gibbs sampler for **BART** is to employ a form of “Bayesian backfitting” (Hastie and Tibshirani, 2000) where the j th tree is fit iteratively, holding all other $m - 1$ trees constant by exposing only the residual response that remains unfitted:

$$\mathbf{R}_j := \mathbf{y} - \sum_{t \neq j} \mathfrak{F}_t(\mathbf{X}). \quad (2)$$

The Gibbs sampler,

$$\begin{aligned} 1 : & \quad \mathfrak{F}_1 \mid \mathbf{R}_{-1}, \sigma^2 \\ 2 : & \quad \mathcal{L}_1 \mid \mathfrak{F}_1, \mathbf{R}_{-1}, \sigma^2 \\ 3 : & \quad \mathfrak{F}_2 \mid \mathbf{R}_{-2}, \sigma^2 \\ 4 : & \quad \mathcal{L}_2 \mid \mathfrak{F}_2, \mathbf{R}_{-2}, \sigma^2 \\ & \quad \vdots \\ 2m - 1 : & \quad \mathfrak{F}_m \mid \mathbf{R}_{-m}, \sigma^2 \\ 2m : & \quad \mathcal{L}_m \mid \mathfrak{F}_m, \mathbf{R}_{-m}, \sigma^2 \\ 2m + 1 : & \quad \sigma^2 \mid \mathfrak{F}_1, \mathcal{L}_1, \dots, \mathfrak{F}_m, \mathcal{L}_m, \mathbf{E}, \end{aligned} \quad (3)$$

proceeds by first proposing a change to the first tree’s structure \mathfrak{F} which are accepted or rejected via a Metropolis-Hastings step (Hastings, 1970). Note that sampling from the

posterior of the tree structure does not depend on the leaf parameters, as they can be analytically margined out of the computation (see Appendix A.1). Given the tree structure, samples from the posterior of the b leaf parameters $\mathcal{E}_1 := \{\mu_1, \dots, \mu_b\}$ are then drawn. This procedure proceeds iteratively for each tree, using the updated set of partial residuals \mathbf{R}_j . Finally, conditional on the updated set of tree structures and leaf parameters, a draw from the posterior of σ^2 is made based on the full model residuals $\mathbf{E} := \mathbf{y} - \sum_{t=1}^m \mathbb{1}_t^{\mathcal{E}}(\mathbf{X})$.

Within a given terminal node, since both the prior and likelihood are normally distributed, the posterior of each of the leaf parameters in \mathcal{E} is conjugate normal with its mean being a weighted combination of the likelihood and prior parameters (lines 2, 4, \dots , $2m$ in equation set 3). Due to the normal-inverse-gamma conjugacy, the posterior of σ^2 is inverse gamma as well (line $2m + 1$ in equation set 3). The complete expressions for these posteriors can be found in Gelman et al. (2004).

Lines 1, 3, \dots , $2m - 1$ in equation set 3 rely on Metropolis-Hastings draws from the posterior of the tree distributions. These involve introducing small perturbations to the tree structure: growing a terminal node by adding two daughter nodes, pruning two daughter nodes (rendering their parent node terminal), or changing a split rule. We denote these possible alterations as: GROW, PRUNE, and CHANGE.¹ The mathematics associated with the Metropolis-Hastings step is not difficult but is tedious, and we refer the interested reader to Appendix A for the explicit calculations. Once again, over many Gibbs samples, trees can dynamically morph their structure in an effort to capture the fit left currently unexplained.

Pratola et al. (2013) argue that a CHANGE step is unnecessary for sufficient mixing of the Gibbs sampler. While we too observed this to be true for estimates of the posterior means, we found that omitting CHANGE could have negative impacts on the variable inclusion proportions (the feature introduced in Section 4.6). As a result, we implement a modified CHANGE where we only propose new splits for nodes that are singly internal: both children nodes are terminal nodes (details are given in Appendix A.3).

All $2m + 1$ steps represent a *single* Gibbs iteration. We’ve found that generally no more than 1,000 iterations are needed as “burn-in” (see Section 4.4 for convergence diagnostics). Then an additional 1,000 iterations is usually sufficient to serve as draws from the posterior for $f(\mathbf{x})$. A single predicted value $\hat{f}(\mathbf{x})$ can be obtained by taking the average of the posterior values and a quantile estimate can be obtained by computing the appropriate quantile of the posterior values. Additional features of the posterior distribution will be discussed in Section 4.

2.3 BART for Classification

BART can easily be modified to handle classification problems for categorical response variables. In Chipman et al. (2010), only binary outcomes were explored but recent work has extended BART to the multiclass problem (Kindo et al., 2013). Our implementation currently handles binary classification problems and we view implementing the extension to multiclass outcomes as fruitful future work.

¹In the original formulation, Chipman et al. (2010) include an additional alteration called SWAP. Due to the complexity of bookkeeping associated with this alteration, we do not implement it.

For the binary classification problem (coded with outcomes “0” and “1”), we assume a probit model,

$$\mathbb{P}(Y = 1 \mid \mathbf{X}) = \Phi \left(\mathfrak{F}_1^{\text{leaf}}(\mathbf{X}) + \mathfrak{F}_2^{\text{leaf}}(\mathbf{X}) + \dots + \mathfrak{F}_m^{\text{leaf}}(\mathbf{X}) \right),$$

where Φ denotes the cumulative density function of the standard normal distribution. In this formulation, the sum-of-trees model serves as an estimate of the conditional probit at \mathbf{x} which can be easily transformed into a conditional probability estimate of $Y = 1$.

In the classification setting, the prior on σ^2 is not needed as the model assumes $\sigma^2 = 1$. The prior on the tree structure remains the same as in the regression setting and a few minor modifications are required for the prior on the leaf parameters.

Sampling from the posterior distribution is again obtained via Gibbs sampling with a Metropolis-Hastings step outlined in Section 2.2. Following the data augmentation approach of (Albert and Chib, 1993) an additional vector of latent variables \mathbf{Z} is introduced into the Gibbs sampler. Then, a new step is created in the Gibbs sampler where draws of $\mathbf{Z} \mid \mathbf{y}$ are obtained by conditioning on the sum-of-trees model:

$$\begin{aligned} Z_i \mid y_i = 1 &\sim \max \left\{ N \left(\sum_t \mathfrak{F}_t^{\text{leaf}}(\mathbf{X}), 1 \right), 0 \right\} \text{ and} \\ Z_i \mid y_i = 0 &\sim \min \left\{ N \left(\sum_t \mathfrak{F}_t^{\text{leaf}}(\mathbf{X}), 1 \right), 0 \right\}. \end{aligned}$$

Next, \mathbf{Z} is used as the response vector instead of \mathbf{y} in all steps of Equation 3.

Upon obtaining a sufficient number of samples from the posterior, inferences can be made using the the posterior distribution of conditional probabilities and classification can be undertaken by applying a threshold to the to the means (or another summary) of these posterior probabilities. The relevant classification features of `bartMachine` are discussed in Section 5.

3 The `bartMachine` package

The package `bartMachine` provides a novel implementation of Bayesian Additive Regression Trees in R. The algorithm is substantially faster than the current R package `BayesTree` and our implementation is parallelized at the Gibbs sample level during prediction. Additionally, the interface with `rJava` allows for the entire posterior distribution of tree ensembles to persist throughout the R session, allowing for prediction and other calls to the trees without having to re-run the Gibbs sampler (a limitation in the current implementation). The model object cannot persist across sessions (using R’s `save` command for instance) and we view the addition of this feature as future work. Since our implementation is different from `BayesTree`, we provide a predictive accuracy bakeoff on different datasets in Appendix B which illustrates that the two are about equal.

3.1 Speed Improvements and Parallelization

We make a number of significant speed improvements over the original implementation.


First, `bartMachine` is fully parallelized (with the number of cores customizable) during model creation, prediction, and many of the other features. During model creation, we chose to parallelize by creating one independent Gibbs chain per core. Thus, if we have 500 burn-in samples and 1,000 post burn-in samples and four cores, each core would sample 750 samples: 500 for a burn-in and 250 post burn-in samples. The final model will aggregate the 250 post burn-in samples for the four cores yielding the desired 1,000 total burn-in samples. This has the drawback of effectively running the burn-in serially, but has the added benefit of reducing auto-correlation of the sum-of-trees samples in the posterior samples since the chains are independent which may provide greater predictive performance. Further parallelization at the level of likelihood calculations we regard as fruitful future work. Parallelization for prediction and other features scale linearly in the number of cores.

Additionally, we take advantage of a number of additional computational shortcuts:

1. Computing the unfitted responses for each tree (Equation 2) can be accomplished by keeping a running vector and making entry-wise updates as the Gibbs sampler (Equation 3) progresses from step 1 to $2m$. Additionally, during the σ^2 step $2m + 1$, the residuals do not have to be computed by dropping the data down all the trees.
2. Each node caches its acceptable variables for split rules and the acceptable unique split values so they do not need to be calculated at each tree sampling step. This speed enhancement, which we call *memcache* comes at the expense of memory and may cause issues for large data sets. Thus we include a toggle in our implementation defaulted to “on.”
3. Careful calculations in Appendix A eliminate many unnecessary computations. For instance, the likelihood ratios are only functions of the squared sum of responses and no longer require computing the sum of the responses squared.

Figure 1 displays model creation speeds for different values of n on a linear model with $p = 20$, normally distributed covariates, $\beta_1, \dots, \beta_{20} \stackrel{iid}{\sim} U(-1, 1)$, and standard normal noise. Note that we do not vary p as it was already shown in Chipman et al. (2010) that BART’s computation time is largely unaffected by the dimensionality of the problem (relative to the influence of sample size). We include results for BART using `BayesTree` (Chipman et al., 2010), `bartMachine` with one and four cores, the *memcache* option on and off, as well as four cores, *memcache* off and computation of in-sample statistics off (all with $m = 50$ trees). We also include Random Forests via the package `randomForest` (Liaw and Wiener, 2002) with its default settings.

We first note that Figure 1a demonstrates that the `bartMachine` model creation runtime is approximately linear in n . There is about a 30% speed-up when using four cores instead of one. The *memcache* enhancement should be turned off only with sample sizes larger than $n = 20,000$. Noteworthy is the 50% reduction in time of constructing the model when not computing in-sample statistics. In-sample statistics are computed by default because the user generally wishes to see them. Also, for the purposes of this comparison, `BayesTree`



(a) Large sample sizes

(b) Small sample sizes

Figure 1: Model creation times as a function of sample size for a number of settings of `bartMachine`, `BayesTree` and `RandomForests`. Simulations were run on a quad-core 3.4GHz Intel i5 desktop with 24GB of RAM running the Windows 7 64bit operating system.

models compute the in-sample statistics by necessity since the trees are not saved. The `randomForest` implementation becomes slower just after $n = 1,000$ due to its reliance on a greedy exhaustive search at each node.

Many users might be interested in low sample sizes ($n \leq 2,000$) whose model creation times are illustrated in Figure 1b. We observe the *memcache* enhancement provides about a 10% speed improvement. Thus, if memory is an issue, it can be turned off without too much of a performance degradation.

3.2 Missing Data in BART

Our package implements a native method for incorporating missing data into both model creation and future prediction with test data. The details are given in Kapelner and Bleich (2013) but we provide a brief summary here.

There are a number of ways to incorporate missingness into tree-based methods (see Ding and Simonoff, 2010 for a review). The method implemented here is known as “Missing Incorporated in Attributes” (MIA, Twala et al., 2008, section 2) which natively incorporates missingness by augmenting the nodes’ splitting rules to (1) also handle sorting the missing data to the left or right and (2) use missingness *itself* as a variable to be considered in a splitting rule. We summarize these new splitting rules as it is implemented within our package in Algorithm 1.

Implementing MIA into the BART procedure is straightforward. These new splitting rules are sampled uniformly during the GROW or CHANGE steps. For example, a splitting rule

might be “ $\mathbf{x}_j < c$ or \mathbf{x}_j is missing.” To account for splitting on missingness itself, we create dummy vectors of length n for each of the p attributes, denoted $\mathbf{M}_1, \dots, \mathbf{M}_p$, which assume the value 1 when the entry is missing and 0 when the entry is present. The original training matrix is then augmented with these dummies, giving the opportunity to select missingness *itself* when choosing a new splitting rule during the grow or change steps. Note that this can increase the number of predictors by up to a factor of 2. We illustrate the building a BART model with missing data in Section 4.9. As described in Chipman et al. (2010, Section 6), BART’s runtime increases negligibly in the number of covariates and this has been our experience using the augmented training matrix.

Algorithm 1 The MIA choices for all attributes $j \in \{1, \dots, p\}$ and all split points x_{ij}^* where $i \in \{1, \dots, n\}$ during a GROW or CHANGE step in `bartMachine`.

- 1: If x_{ij} is missing, send it \leftarrow ; if it is present and $x_{ij} \leq x_{ij}^*$, send it \leftarrow , otherwise \rightarrow .
 - 2: If x_{ij} is missing, send it \rightarrow ; if it is present and $x_{ij} \leq x_{ij}^*$, send it \leftarrow , otherwise \rightarrow .
 - 3: If x_{ij} is missing, send it \leftarrow ; if it is present, send it \rightarrow .
-

3.3 Principled Variable Selection

Our package also implements the variable selection procedures developed in Bleich et al. (2013), which is best applied to data problems where the number of covariates influencing the response is small relative to the total number of covariates. We give a brief summary of the procedures here.

In order to select variables, we make use of the “variable inclusion proportions,” the proportion of times each predictor is chosen as a splitting rule divided by the total number of splitting rules appearing in the model (see Section 4.6 for more details). The variable selection procedure can be outlined as follows:

1. Compute the model’s variable inclusion proportions.
2. Permute the response vector, regenerate the model, and compute the “null” variable inclusion proportions. Repeat this a number of times to create a “null” permutation distribution.
3. Three selection rules can be used depending on the desired stringency of selection:
 - (a) Local Threshold: Include a predictor \mathbf{x}_k if its variable inclusion proportion exceeds the $1 - \alpha$ quantile of its own null distribution.
 - (b) Global Max Threshold: Include a predictor \mathbf{x}_k if its variable inclusion proportion exceeds the $1 - \alpha$ quantile of the distribution of the maximum of the null variable inclusion proportions from each permutation of the response.
 - (c) Global SE Threshold: Select \mathbf{x}_k if its variable inclusion proportion exceeds a threshold based from its own null distribution mean and SD with a global multiplier shared by all predictors.

The Local procedure is the least stringent in terms of selection and the Global Max procedure the most. The Global SE procedure is a compromise, but behaves more similarly to the Global Max. Bleich et al. (2013) demonstrate that the best procedure depends on the underlying sparsity of the problem, which is often unknown. Therefore, the authors include an additional procedure that chooses the best of these thresholds via cross-validation and this method is also implemented in `bartMachine`. Examples of these procedures for variable selection are provided in Section 4.10.

4 Regression Features

We illustrate the package features by using both real and simulated data, focusing first on regression problems.

4.1 Computing parameters

We first set some computing parameters. We allow up to 5GB of RAM for the Java heap (although we never used more than 1GB during this paper’s exploration)² and we set the number of computing cores available for use to 4.

```
> set_bart_machine_num_cores(4)
> init_java_for_bart_machine_with_mem_in_mb(5000)
```

The following Sections (4.2 – 4.10) use a dataset obtained from UCI (Bache and Lichman, 2013). The $n = 201$ observations are automobiles and the goal is to predict their price from 25 features (15 continuous and 10 nominal), first explored by Kibler et al. (1989).³ This dataset also contains missing data.

4.2 Model Building

We are now are ready to construct a `bartMachine` model. The default hyperparameters generally follow the recommendations of Chipman et al. (2010) and provide a ready-to-use algorithm for many data problems. Our hyperparameter settings are $m = 50$,⁴ $\alpha = 0.95$, $\beta = 2$, $k = 2$, $q = 0.9$, $\nu = 3$, probabilities of the GROW / PRUNE / CHANGE steps is 39% / 39% / 44%. We set the number of burn-in Gibbs samples to 250 and number of post-burn-in samples to 1,000. We default the missing data feature to be off. We default

²Note that the maximum amount of memory can be set only *once* at the beginning of the R session (a limitation of `rJava` since only one Java Virtual Machine can be initiated per session), but the number of cores can be respecified at any time.

³We first preprocess the data. We first drop one of the nominal predictors (car company) due to too many categories (22). We then coerce two of the of the nominal predictors to be continuous. Further, the response variable, price, was logged to reduce right skew in its distribution.

⁴In contrast to Chipman et al. (2010), we recommend this default as a good starting point rather than $m = 200$ due to our experience experimenting with the “RMSE by number of trees” feature found in later in this section. Performance is often similar and computational time and memory requirements are dramatically reduced.

the covariates to be equally important *a priori*. Other parameters and their defaults can be found in the package's online manual. Below is a default `bartMachine` model.

```
> bart_machine = build_bart_machine(X, y)
Building BART for regression ... evaluating in sample data...done
```

If one wishes to see the iterations of the Gibbs sampler of Equation 3, the flag `verbose` can be set to "TRUE". One can see more debug information from the Java program by setting the flag `debug_log` to true and the program will print to `unnamed.log` in the current working directory. We now inspect the model object to query its in-sample performance and to be reminded of the input data and model hyperparameters.

```
> bart_machine
Bart Machine v1.0b for regression

training data n = 160 and p = 46
built in 1 secs on 4 cores, 50 trees, 250 burn in and 1000 post. samples

sigsq est for y beforehand: 0.014
avg sigsq estimate after burn-in: 0.00886

in-sample statistics:
  L1 = 9.03
  L2 = 0.8
  rmse = 0.07
  Pseudo-Rsq = 0.9741
p-val for shapiro-wilk test of normality of residuals: 0.01785
p-val for zero-mean noise: 0.97692
```

Figure 2: The summary for the default `bartMachine` model built with the automobile data

Since the response was considered continuous, we employ BART for regression. The dimensions of the design matrix are given. Note that we dropped 45 observations that contained missing data (which we will retain in Section 4.9). We then have a recording of the MSE for the OLS model and our average estimate of σ_e^2 . We are then given in-sample statistics on error. Pseudo- R^2 is calculated via $1 - SSE/SST$. Also provided are outputs from tests of the error distribution being mean centered and normal. In this case, we cannot conclude normality of the residuals using the Shapiro-Wilk test.

We can also obtain out-of-sample statistics to assess level of overfitting by using k-fold cross validation. Using 10 folds we find:

```
> k_fold_cv(X, y, k_folds = 10)
$L1_err      $L2_err      $rmse      $PseudoRsq
[1] 22.63155   [1] 5.202831   [1] 0.1803266 [1] 0.831917
```

The Pseudo- R^2 being lower out-of-sample versus in-sample suggests evidence that BART is slightly overfitting (note that the training sample during cross-validation is 10% smaller).

It may also be of interest how the number of trees m affects performance. One can examine how out-of-sample predictions vary by the number of trees via

```
> rmse_by_num_trees(bart_machine, num_replicates = 20)
```

and the output is shown in Figure 3.




Figure 3: Out-of-sample predictive performance by number of trees

It seems that increasing $m > 50$ does not result in any substantial increase in performance. We can now try to build a better `bartMachine` by grid-searching over a set of hyperparameter combinations, including m (for more details, see BART-cv in Chipman et al., 2010). The default grid search is small and it can be customized by the user.

```
> bart_machine_cv = build_bart_machine_cv(X, y)
...
BART CV win: k: 2 nu, q: 10, 0.75 m: 200
```

This function returns the “winning” model, which is the one with lowest out-of-sample RMSE over a 5-fold cross-validation. Here, the cross-validated `bartMachine` model has slightly better in-sample performance ($L1 = 8.18$, $L2 = 0.68$ and $\text{Pseudo-}R^2 = 0.978$) as well as slightly better out-of-sample performance:

```
> k_fold_cv(X, y, k_folds = 10, k = 2, nu = 3, q = 0.9, num_trees = 200)
$L1_err      $L2_err      $rmse      $PseudoRsq
[1] 21.21557   [1] 4.517916   [1] 0.1680386 [1] 0.8540439
```

Predictions are handled with the `predict` function:

```
predict(bart_machine_cv, X[1 : 14, ])  
  9.479963  9.775766  9.799110 10.050041  9.659138  9.697902  9.873622  
  9.931972  8.550436  8.688874  8.794351  8.647986  8.690300  9.029066
```

We also include a convenience method `bart_predict_for_test_data` that will predict and return out-of-sample error metrics when the test outcomes are known.

4.3 Model Destroying

As noted in the introduction to Section 3, `bartMachine` objects persist in Java for the entirety of an R session. These model objects can use a substantial amount of RAM and it is prudent to release this memory when it is no longer needed. Simply removing a `bartMachine` in R does *not* destroy the Java object and release the RAM (resulting in a memory leak).

Therefore, we provide a utility function `destroy_bart_machine` that cleans up the Java object. This function should be called when a `bartMachine` object is no longer needed and *before* removing or overwriting the R variable. Since we no longer are using the original `bart_machine` object, we now release its memory via:

```
destroy_bart_machine(bart_machine)
```

4.4 Assumption Checking

BART relies on a number of assumptions for obtaining valid Bayesian inference. We include features that assess the plausibility of these assumptions.

Checking the mean-centeredness of the noise is addressed in the summary output of Figure 2 and is simply a one-sample *t*-test of the average residual value against a null hypothesis of true mean zero.

We next assess both normality and heteroskedasticity using the following command:

```
> check_bart_error_assumptions(bart_machine_cv)
```

This will display a window similar to Figure 4 which contains a QQ-plot (to assess normality) as well as a residual-by-predicted plot (to assess homoskedasticity). It appears that the errors are most likely normal and homoskedastic.

In addition to the model assumptions, BART requires convergence of its Gibbs sampler. Figure 5 displays four types of convergence diagnostics.

4.5 Credible Intervals and Prediction Intervals

An advantage of BART over similar machine learning algorithms is that if the priors and model assumptions are believed, the Bayesian probability model and corresponding burned-in Gibbs samples provide the approximate posterior distribution of $f(\mathbf{x})$. Thus, one can compute uncertainty estimates via quantiles of the posterior samples. These provide Bayesian “credible intervals” which are intervals for the conditional expectation function, $\mathbb{E}[\mathbf{y} \mid \mathbf{X}]$.




Figure 4: Test of normality of errors using QQ-plot and the Shapiro-Wilk test (top), residual plot to assess heteroskedasticity (bottom).




Figure 5: Convergence diagnostics for the cross-validated `bartMachine` model. Top left: σ^2 by Gibbs sample. Samples to the left of the first vertical grey line are burn-in from the first computing core's Gibbs sample chain. The four subsequent plots separated by grey lines are the post-burn-in samples from each of the four computing cores employed during model construction. Top right: percent acceptance of Metropolis-Hastings proposals across the m trees where each point plots one iteration. Points before the grey vertical line illustrate burn-in samples and points after illustrate post burn-in. Each computing core is colored differently. Bottom left: average number of leaves across the m trees by sample (post burn-in only where computing cores separated by vertical grey lines). Bottom right: average tree depth across the m trees by sample (post burn-in only where computing cores separated by vertical grey lines).

Another useful uncertainty interval can be computed for individual predictions by combining uncertainty from the conditional expectation function with the systematic, homoskedastic normal noise produced by \mathcal{E} . Since we have draws from the posterior of the conditional expectation distribution and concomitant draws from the posterior of the variance distribution, we can simulate the distribution of the response by drawing many realizations from $\mathcal{N}(f(\mathbf{x})_n, \sigma_n^2)$ for each the post-burn-in samples and aggregating them together. The prediction interval is then provided by the appropriate quantiles of the posterior samples.

Below is an example of how both types of intervals are computed in the package (for the 100th observation of the training data):

```
> calc_credible_intervals(bart_machine_cv, X[100, ], ci_conf = 0.95)
      ci_lower_bd ci_upper_bd
[1,]    8.725202    8.971687
> calc_prediction_intervals(bart_machine_cv, X[100, ], pi_conf = 0.95)
      pi_lower_bd pi_upper_bd
[1,]    8.631243    9.06353
```

Note that the prediction intervals are wider than the credible intervals because they reflect the uncertainty from the error term.

We can then plot these intervals in sample:

```
> plot_y_vs_yhat(bart_machine_cv, credible_intervals = TRUE)
> plot_y_vs_yhat(bart_machine_cv, prediction_intervals = TRUE)
```


Figure 6a shows how our prediction fared against the original response (in-sample) with 95% credible intervals. Figure 6b shows the same prediction versus the original response plot now with 95% prediction intervals.

4.6 Variable Importance


After a BART model is built, it is natural to ask the question: which variables are most important? This is assessed by examining the splitting rules in the m trees across the post burn-in gibbs samples which are known as “inclusion proportions” (Chipman et al., 2010). The inclusion proportion for any given predictor represents the proportion of times that variable is chosen as a splitting rule out of all splitting rules among the posterior draws of the sum-of-trees model. Figure 7 illustrates the inclusion proportions for all variables obtained via:

```
> investigate_var_importance(bart_machine_cv, num_replicates_for_avg = 20)
```

Actual selection of variables *significantly* affecting the response is addressed conceptually in Section 3.3 and examples are provided in Section 4.10.



(a) Segments illustrate credible intervals



(b) Segments illustrate prediction intervals

Figure 6: Fitted versus actual response values for the automobile dataset. Segments are 95% credible intervals (a) or 95% prediction intervals (b). Green dots indicate the true response is within the stated interval and red dots indicate otherwise. Note that the percent coverage in (a) is not expected to be 95% because the response includes a noise term.




Figure 7: Average variable inclusion proportions in the cross-validated `bartMachine` model for the automobile data averaged over 100 model constructions to obtain stable estimates across many posterior modes in the sum-of-trees distribution (as recommended in Bleich et al., 2013). The segments atop the bars represent 95% confidence intervals. The eight predictors with inclusion proportions of zero feature identically one value (after missing data was dropped).

4.7 Variable Effects

It is also natural to ask: does \mathbf{x}_j affect the response, controlling for other variables in the model? This is roughly analogous to the t -test in ordinary least squares regression of no linear effect of \mathbf{x}_j on \mathbf{y} while controlling for \mathbf{x}_{-j} . The null hypothesis here is the same but the linearity constraint is relaxed. To test this, we employ a permutation approach where we record the observed Pseudo- R^2 from the `bartMachine` model built with the original data. Then we permute the \mathbf{x}_j th column, thereby destroying any relationship between \mathbf{x}_j and \mathbf{y} , construct a new duplicate `bartMachine` model from this permuted design matrix and record a “null” Pseudo- R^2 . We then repeat this process to obtain a null distribution of Pseudo- R^2 ’s. Since the alternative hypothesis is that \mathbf{x}_j has an effect on \mathbf{y} in terms of predictive power, our p_{val} is the proportion of null Pseudo- R^2 ’s greater than the observed Pseudo- R^2 , making our procedure a natural one-sided test. Note, however, that this test is conditional on the `BART` model and its selected priors being true (similar to the assumptions of the linear model).

If we wish to test if a set $A \subset \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ affect the response after controlling for other variables, we repeat the procedure outlined in the above paragraph by permuting the columns of A in every null sample. This is roughly analogous to the partial F -test in ordinary least squares regression.

If we wish to test if *any* of the covariates matter in predicting \mathbf{y} , roughly analogous to the omnibus F -test in ordinary least squares regression, we simply permute \mathbf{y} during the null sampling. This procedure breaks the relationship between the response and the predictors but does not alter the existing associations between predictors.

At $\alpha = 0.05$, Figure 8a demonstrates an insignificant effect of the variable `width` of car on price. Even though `width` is putatively the “most important” variable as measured

by proportions of splits in the posterior sum-of-trees model (Figure 7), note that this is largely an easy prediction problem with many collinear predictors. Figure 8b shows the results of a test of the putatively most important categorical variable, `body style` (which involves permuting the categories, then dummifying the levels to preserve the structure of the variable). We find a marginally significant effect ($p = 0.0495$). A test of the top ten most important variables is convincingly significant (Figure 8c). For the omnibus test, Figure 8d illustrates an extremely statistically significant result, as would be expected. The code to run these tests is shown below (output suppressed).

```
> cov_importance_test(bart_machine_cv, covariates = c("width"))
> cov_importance_test(bart_machine_cv, covariates = c("body_style"))
> cov_importance_test(bart_machine_cv, covariates = c("width",
  "curb_weight", "city_mpg", "length", "horsepower", "body_style",
  "engine_size", "highway_mpg", "peak_rpm", "normalized_losses"))
> cov_importance_test(bart_machine_cv)
```




Figure 8: Tests of covariate importance conditional on the cross-validated `bartMachine` model. All tests performed with 100 null samples.

4.8 Partial Dependence

A data analyst may also be interested in understanding how x_j affects the response on average, after controlling for other predictors. This can be examined using Friedman (2001)'s Partial Dependence Function (PDP),

$$f_j(\mathbf{x}_j) = \mathbb{E}_{\mathbf{x}_{-j}} [f(\mathbf{x}_j, \mathbf{x}_{-j})] = \int f(\mathbf{x}_j, \mathbf{x}_{-j}) dP(\mathbf{x}_{-j}), \quad (4)$$

where \mathbf{x}_{-j} denotes all variables except \mathbf{x}_j . The PDP of predictor \mathbf{x}_j gives the average value of f when \mathbf{x}_j is fixed and \mathbf{x}_{-j} varies over its marginal distribution, $dP(\mathbf{x}_{-j})$. As neither the true model f nor the distribution of the predictors $dP(\mathbf{x}_{-j})$ are known, we estimate Equation 4 by computing

$$\hat{f}_j(\mathbf{x}_j) = \frac{1}{n} \sum_{i=1}^n \hat{f}(\mathbf{x}_j, \mathbf{x}_{-j,i}) \quad (5)$$

where n is the number of observations in the training data and \hat{f} denotes the `bartMachine` model. Since BART provides an estimated posterior distribution, we can plot credible bands for the PDP function. In Equation 5, the \hat{f} can be replaced with a function that calculates the q th quantile of the post-burned-in Gibbs samples for \hat{y} . Figure 9a plots the PDP along with the 2.5%ile and the 97.5%ile for the variable `horsepower`. By varying over most of the range of `horsepower`, the price is predicted to increase by about \$1000. Figure 9b plots the PDP along with the 2.5%ile and the 97.5%ile for the variable `stroke`. This predictor seemed to be relatively unimportant according to Figure 7 and the PDP confirms this, with a very small, yet nonlinear average partial effect. The code for both plots is below.

```
> pd_plot(bart_machine_cv, j = "horsepower")
> pd_plot(bart_machine_cv, j = "stroke")
```

4.9 Incorporating Missing Data

The procedure for incorporating missing data was introduced in Section 3.2. We now build a `bartMachine` model using this procedure below:

```
> bart_machine = build_bart_machine(X, y, use_missing_data = TRUE,
  use_missing_data_dummies_as_covars = TRUE)
> bart_machine
Bart Machine v1.0b for regression

Missing data feature ON
training data n = 201 and p = 50
built in 1.3 secs on 1 core, 50 trees, 250 burn-in and 1000 post. samples

sigsq est for y beforehand: 0.016
avg sigsq estimate after burn-in: 0.01055

in-sample statistics:
L1 = 12.77
```




Figure 9: PDPs plotted in black and 95% credible intervals plotted in blue for variables in the automobile dataset. Points plotted are at the 5%ile, 10%ile, 20%ile, ..., 90%ile and 95%ile of the values of the predictor. Lines plotted between the points approximate the PDP by linear interpolation.

```

L2 = 1.28
rmse = 0.08
Pseudo-Rsq = 0.9746
p-val for shapiro-wilk test of normality of residuals: 0.6638
p-val for zero-mean noise: 0.95693

```

Note that we are now using the complete data set including the 41 observations where there were missing features. Also note that p has now increased from 46 to 51. The five new predictors are dummy variables which indicate missingness constructed from the predictors which exhibited missingness (due to the `use_missing_data_dummies_as_covars` parameter being set to true). These variables are important if missingness *itself* shifts the response, as was the case in models explored in Kapelner and Bleich (2013). One way to test if this type of missingness is important is to run a covariate test (Section 4.7) on these new 5 covariates:

```

> cov_importance_test(bart_machine, covariates = c("M_normalized_losses",
  "M_bore", "M_stroke", "M_horsepower", "M_peak_rpm"))
BART test for importance of 5 covariates....p_val = 0.673

```

We cannot conclude that these variables have an effect on the response (which seems reasonable given that only 5 predictors featured missingness among only 40 observations). This suggests that it may be appropriate to shrink the model by leaving out these missing dummies, but still allow for missingness to be incorporated into the split rules. We drop these predictions and rebuild the model for the rest of this exposition.

The procedure of Section 3.2 also natively incorporates missing data during prediction. Missingness will yield larger credible intervals. In the example below, we suppose that the `curb_weight` and `symboling` values were unavailable for 20th automobile.

```

> x_star = X[20, ]
> calc_credible_intervals(bart_machine, x_star, ci_conf = 0.95)
      ci_lower_bd ci_upper_bd
[1,]      8.650093      8.824515
> x_star[c("curb_weight", "symboling")] = NA
> calc_credible_intervals(bart_machine, x_star, ci_conf = 0.95)
      ci_lower_bd ci_upper_bd
[1,]      8.622582      8.978313

```

4.10 Variable Selection

In this section we implement a principled variable selection procedure introduced in Section 3.3. The following code will select variables based on the three thresholds and also displays the plot in Figure 10.⁵

```

> vs = var_selection_by_permute_response_three_methods(bart_machine,
      bottom_margin = 10, num_permute_samples = 10)
> vs$important_vars_local_names
"curb_weight" "city_mpg" "engine_size" "horsepower"
"length"      "width"      "num_cylinders" "body_style_convertible"
"wheel_base"  "peak_rpm" "highway_mpg"  "wheel_drive_fwd"
> vs$important_vars_global_max_names
"curb_weight" "city_mpg" "engine_size" "horsepower" "length"
> vs$important_vars_global_se_names
"curb_weight" "city_mpg" "engine_size" "horsepower" "length"
"width"      "num_cylinders" "wheel_base" "wheel_drive_fwd"

```

Usually, “Global Max” and “Global SE” perform similarly, as they are both more stringent in selection. However, in many situations it will not be clear to the data analyst which threshold is most appropriate. The “best” procedure can be chosen via cross-validation on out-of-sample RMSE as follows:

```

var_selection_by_permute_response_cv(bart_machine)
$best_method
[1] "important_vars_local_names"

$important_vars_cv
 [1] "body_style_convertible" "city_mpg"          "curb_weight"
 [4] "engine_size"           "engine_type_ohc"   "horsepower"
 [7] "length"                "num_cylinders"     "peak_rpm"
[10] "wheel_base"            "wheel_drive_fwd"   "wheel_drive_rwd"
[13] "width"

```

⁵By default, variable selection is performed individually on dummy variables for a factor. The variable selection procedures return the permutation distribution and an aggregation of the dummy variables’ inclusion proportions can allow for variable selection to be performed on an entire factor.




Figure 10: Visualization of the three variable selection procedures outlined in Section 3.3 with $\alpha = 0.05$. The top plot illustrates the “Local” procedure. The green lines are the threshold levels determined from the permutation distributions that must be exceeded for a variable to be selected. The plotted points are the variable inclusion proportions for the observed data (averaged over five duplicate `bartMachine` models). If the observed value is higher than the green bar, the variable is included and is displayed as a solid dot; if not, it is not included and it is displayed as an open dot. The bottom plot illustrates both the “Global SE” and “Global Max” thresholds. The red line is the cutoff for “Global Max” and variables pass this threshold are displayed as solid dots. The blue lines represent the thresholds for the “Global SE” procedure. Variables that exceed this cutoff but not the “Global Max” threshold are displayed as asterisks. Open dots exceed neither threshold.

On this dataset, the “best” approach (as defined by out-of-sample prediction error) is the “Local” procedure.

The following sections (4.11 and 4.12) demonstrate additional features using Friedman (1991)’s function:

$$\mathbf{y} = 10\sin(\pi\mathbf{x}_1\mathbf{x}_2) + 20(\mathbf{x}_3 - .5)^2 + 10\mathbf{x}_4 + 5\mathbf{x}_5 + \mathcal{E}, \quad \mathcal{E} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}). \quad (6)$$

4.11 Informed Prior Information on Covariates

Bleich et al. (2013) propose a method for incorporating informed prior information about the predictors into BART. This can be achieved by modifying the prior on the splitting rules as well as the corresponding calculations in the Metropolis-Hastings step. In particular, covariates believed to influence the response can be proposed more often as candidates for splitting rules. Useful prior information can aid in both variable selection and prediction tasks. We illustrate the impact of a correctly informed prior in the context of the Friedman function (Equation 6). We include the 5 predictors which influence the response as well as 95 that do not.

All that is required is a specification of relative weights for each predictor. These are converted internally to probabilities. We assign 5 times the weight to the 5 true covariates of the model relative to the 95 useless covariates.

```
> prior = c(rep(5, times = 5), rep(1, times = 95))
```

We now sample 500 observations from the Friedman function and construct a default `bartMachine` model as well as a `bartMachine` model with the informed prior and compare their performance on a test set of another 500 observations.

```
> bart_machine = build_bart_machine(X, y)
> bart_machine_informed = build_bart_machine(X, y, cov_prior_vec = prior)

> bart_predict_for_test_data(bart_machine, Xtest, ytest)$rmse
[1] 1.661159
> bart_predict_for_test_data(bart_machine_informed, Xtest, ytest)$rmse
[1] 1.232925
```

There is a substantial improvement in out-of-sample predictive performance when a properly informed prior is used.

Note that we recommend use of the prior vector to down-weight the indicator variables that result from dummifying factors so that the total set of dummy variables has the same weight as a continuous covariate.

4.12 Interaction Effect Detection

In Section 4.6, we explored using variable inclusion proportions to understand the relative influences of different covariates. A similar procedure can be carried out for examining interaction effects within a BART model. This question was initially explored in Damien et al. (2013) where an interaction was considered to exist between two variables if they both appeared in at least one splitting rule in a given tree. We refine the definition of an interaction as follows.

We first begin with a $p \times p$ matrix of zeroes. Within a given tree, for each split rule variable j , we look at all split rule variables of child nodes, k , and we increment the j, k element of the matrix. Hence variables are considered to interact in a given tree *only if* they appear together in a contiguous downward path from the root node to a terminal node. Note that a variable may interact with itself (when fitting a linear effect, for instance). Since there is no order between the parent and child, we then add the j, k counts together with the k, j counts (if $j \neq k$). Summing across trees and Gibbs samples gives the total number of interactions for each pair of variables from which relative importance can be assessed.

We demonstrate interaction detection on the Friedman function using 10 covariates using the code below:

```
> interaction_investigator(bart_machine, num_replicates_for_avg = 25,
  num_var_plot = 10, bottom_margin = 5)
```

Shown in Figure 11 are the 10 most important interactions in the model. The illustration is averaged over many model constructions to obtain stable estimates across many posterior modes in the sum-of-trees distribution. Notice that the interaction between x_1 and x_2 dominates all other terms, as BART is correctly capturing the single true interaction effect in Equation 6. Choosing which of these interactions *significantly* affect the response is not addressed in this paper but we believe methods as suggested in Section 3.3 can be applied and we consider this to be fruitful future work.




Figure 11: The top 10 average variable interaction counts (termed “relative importance”) in the default `bartMachine` model for the Friedman function data averaged over 25 model constructions. The segments atop the bars represent 95% confidence intervals.

5 Classification Features

In this section, we highlight the features that differ from the regression case when the response is dichotomous. The illustrative dataset consists of 332 Pima Indians obtained from the UCI repository. Of the 332 subjects, 109 were diagnosed with diabetes, the binary response variable. There are seven continuous predictors which are body metrics such as blood pressure, glucose concentration, etc. and there is no missing data.

Building a `bartMachine` model for classification has the same computing parameters except that q, ν cannot be specified since there is no longer a prior on σ^2 (see Section 2.3). We first build a cross-validated model below.

```
> bart_machine_cv = build_bart_machine_cv(X, y)
... BART CV win: k: 3 m: 50
> bart_machine_cv
Bart Machine v1.0b for classification
```

```
training data n = 332 and p = 7
built in 0.5 secs on 4 cores, 50 trees, 250 burn-in and 1000 post. samples
```

confusion matrix:

	predicted No	predicted Yes	model errors
actual No	211.000	12.00	0.054
actual Yes	41.000	68.00	0.376
use errors	0.163	0.15	0.160

Classification models have an added hyperparameter, `prob_rule_class`, which is the rule for determining if the probability estimate is great enough to be classified into the positive category. We can see above that the `bartMachine` at times predicts “NO” for true “YES” outcomes and we suffer from a 37.6% error rate for this outcome. We can try to mitigate this error by lowering the threshold to increase the number of “YES” labels predicted:

```
> build_bart_machine(X, y, prob_rule_class = 0.3)
Bart Machine v1.0b for classification
```

```
training data n = 332 and p = 7
built in 0.5 secs on 4 cores, 50 trees, 250 burn-in and 1000 post. samples
```

confusion matrix:

	predicted No	predicted Yes	model errors
actual No	178.000	45.000	0.202
actual Yes	12.000	97.000	0.110
use errors	0.063	0.317	0.172

The model error dropped down to 11% for the “YES” class, but this comes at the expense of increasing the error rate for the “NO” class. We encourage the user to cross-validate this rule based on the appropriate objective function for the problem at hand.

We can now check out-of-sample statistics:

```
> oos_stats = k_fold_cv(X, y, k_folds = 10)
> oos_stats$confusion_matrix
```

	predicted No	predicted Yes	model errors
actual No	203.000	20.000	0.090
actual Yes	47.000	62.000	0.431
use errors	0.188	0.244	0.202

Note that it is possible to predict both class labels and probability estimates for given observations:

```
> predict(bart_machine_cv, X[1 : 2, ], type = "prob")
[1] 0.6253160 0.1055975
> predict(bart_machine_cv, X[1 : 2, ], type = "class")
[1] "Yes" "No"
```

When using the covariate tests of Section 4.7, total misclassification error becomes the statistic of interest instead of Pseudo- R^2 . The p value is calculated now as the proportion of null samples with *higher* misclassification error. Figure 12 illustrates the test showing that predictor `age` seems to matter in the prediction of Diabetes, controlling for other predictors.




Figure 12: Test of covariate importance for predictor `age` on whether or not the subject will contract Diabetes.

The partial dependence plots of Section 4.8 are now scaled as probit of the probability estimate. Figure 13 illustrates that as glucose increases, the probability of contracting Diabetes increases linearly on a probit scale.

Credible intervals are implemented for classification `bartMachine` and are displayed on the probit scale. Note that the prediction intervals of Section 4.5 do not exist for classification.




Figure 13: PDP for predictor `glu`. The blue lines are 95% credible intervals.

```
> calc_credible_intervals(bart_machine_cv, X[1 : 2, ])
      ci_lower_bd ci_upper_bd
[1,]  0.34865355  0.8406097
[2,]  0.01686486  0.2673171
```

Other functions work similarly to regression except those that plot the responses and those that explicitly depend on RMSE as an error metric.

6 Discussion

In this article, we have introduced `bartMachine`, a new R package which implements Bayesian Additive Regression Trees. The goal of this package is to provide a fast, extensive and user-friendly implementation accessible to a wide range of data analysts, and increase the visibility of `BART` to a broader statistical audience. One of the beauties of the `BART` model is its underlying probability model which can be altered for use in different model scenarios. We provide organized, well-documented open-source code and we encourage the community to make innovations on this package.

Replication

The code for `bartMachine` is located at <http://github.com/kapelner/bartMachine> under the MIT license. Results, tables, and figures found in this paper can be replicated via the scripts located in the `bart_package_paper` folder within this git repository.

Acknowledgements

We thank Richard Berk, Andreas Buja, Zachary Cohen, Ed George, Alex Goldstein, Shane Jensen, Abba Krieger, and Robert DeRubeis for helpful discussions. We thank Simon Urbanek for his generous help with `rJava`. Adam Kapelner acknowledges support from the National Science Foundation’s Graduate Research Fellowship Program.

References

- JH Albert and S Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679, June 1993.
- K Bache and M Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Gail Blattenberger and Richard Fowles. Avalanche forecasting: Using bayesian additive regression trees (BART). In *Demand for Communications Services—Insights and Perspectives*, pages 211–227. Springer, 2014.
- J Bleich, A Kapelner, ST Jensen, and EI George. Variable selection inference for bayesian additive regression trees. *ArXiv e-prints*, 2013.
- L Breiman. Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3):199–231, 2001.
- HA Chipman, EI George, and RE McCulloch. BART: Bayesian Additive Regressive Trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- P Damien, P Dellaportas, NG Polson, and DA Stephens. *Bayesian Theory and Applications*, pages 455–464. Oxford University Press, first edition, 2013.
- Y Ding and JS Simonoff. An investigation of missing data methods for classification trees applied to binary response data. *The Journal of Machine Learning Research*, 11:131–170, 2010.
- Jehoshua Eliashberg. *Green-lighting Movie Scripts: Revenue Forecasting and Risk Management*. PhD thesis, University of Pennsylvania, 2010.
- JH Friedman. Multivariate adaptive regression splines (with discussion and a rejoinder by the author). *The annals of statistics*, 19:1–67, 1991.
- JH Friedman. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.
- JH Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- A Gelman, JB Carlin, HS Stern, and DB Rubin. *Bayesian Data Analysis*. Chapman & Hall / CRC, second edition, 2004.

- S Geman and D Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- T Hastie and R Tibshirani. Bayesian Backfitting. *Statistical Science*, 15(3):196–213, 2000.
- WK Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- A Kapelner and J Bleich. Prediction with missing data via bayesian additive regression trees. *ArXiv e-prints*, 2013.
- DF Kibler, DW Aha, and MK Albert. Instance-based prediction of real-valued attributes. *Computational Intelligence*, 5:51, 1989.
- B Kindo, H Wang, and E Pe. MBACT - Multiclass Bayesian Additive Classification Trees. pages 1–29, 2013.
- A Liaw and M Wiener. Classification and regression by randomforest. *R News*, 2(3):18–22, 2002.
- MT Pratola, H Chipman, D Higdon, R McCulloch, and W Rust. Parallel bayesian additive regression trees. Technical report, University of Chicago, 2013.
- MA Taddy, RB Gramacy, and NG Polson. Dynamic Trees for Learning and Design. *Journal of the American Statistical Association*, 106(493):109–123, 2011.
- B Twala, M Jones, and DJ Hand. Good methods for coping with missing data in decision trees. *Pattern Recognition Letters*, 29(7):950–956, 2008.
- S Urbanek. *rJava: Low-level R to Java interface*, 2011. R package version 0.9-3 available on CRAN.
- Qing Zhou and Jun S Liu. Extracting sequence features to predict protein–DNA interactions: a comparative study. *Nucleic acids research*, 36(12):4137–4148, 2008.

A Sampling New Trees

This section provides details on the implementation of Equation 3 (steps 1, 3, \dots , $2m - 1$), the Metropolis-Hastings step for sampling new trees. Recall from Section 2.2 that trees can be altered via growing new daughter nodes from an existing terminal node, pruning two terminal nodes such that their parent becomes terminal, or changing the splitting rule in a node.

Below is the Metropolis ratio (Gelman et al., 2004, p.291) where the parameter sampled is the tree and the data is the responses unexplained by other trees denoted by \mathbf{R} . We denote the new, proposal tree with an asterisk and the original tree without the asterisk.

$$r = \frac{\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right) \mathbb{P}\left(\mathfrak{T}_* \mid \mathbf{R}, \sigma^2\right)}{\mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right) \mathbb{P}\left(\mathfrak{T} \mid \mathbf{R}, \sigma^2\right)} \quad (7)$$

We accept a draw from the posterior distribution of trees if a draw from a standard uniform distribution is less than the value of r . Immediately we note that it is difficult (if not impossible) to calculate the posterior probabilities for the trees themselves. Instead, we employ Bayes' Rule,

$$\mathbb{P}\left(\mathfrak{T} \mid \mathbf{R}, \sigma^2\right) = \frac{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}, \sigma^2\right) \mathbb{P}\left(\mathfrak{T} \mid \sigma^2\right)}{\mathbb{P}\left(\mathbf{R} \mid \sigma^2\right)},$$

and plug the result into Equation 7 to obtain:

$$r = \underbrace{\frac{\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right)}{\mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right)}}_{\text{transition ratio}} \times \underbrace{\frac{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}_*, \sigma^2\right)}{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}, \sigma^2\right)}}_{\text{likelihood ratio}} \times \underbrace{\frac{\mathbb{P}\left(\mathfrak{T}_*\right)}{\mathbb{P}\left(\mathfrak{T}\right)}}_{\text{tree structure ratio}}.$$

Note that the probability of the tree structure is independent of σ^2 .

The goal of this section is to explicitly calculate r for all possible tree proposals — GROW, PRUNE and CHANGE. For each proposal, the calculations are organized into separate sections detailing each of the three ratios — transition, likelihood and tree structure. Note that our actual implementation uses the following expressions in log form for numerical accuracy.

A.1 Grow Proposal

Transition Ratio

Transitioning from the original tree to a new tree involves growing two daughter nodes from a current terminal node:

$$\begin{aligned} \mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right) &= \mathbb{P}(\text{GROW}) \mathbb{P}(\text{selecting } \eta \text{ to grow from}) \times \\ &\quad \mathbb{P}(\text{selecting the } j\text{th attribute to split on}) \times \\ &\quad \mathbb{P}(\text{selecting the } i\text{th value to split on}) \\ &= \mathbb{P}(\text{GROW}) \frac{1}{b} \frac{1}{p_{\text{adj}}(\eta)} \frac{1}{n_{j \cdot \text{adj}}(\eta)}. \end{aligned} \quad (8)$$

We chose one of the current b terminal nodes which we denote the η th node, and then we pick an attribute and split point. $p_{\text{adj}}(\eta)$ denotes the number of predictors left available to split on. This can be less than p if certain predictors do not have two or more unique values once the data reaches the η th node. For example, this regularly occurs if a dummy variable

was split on in some node higher up in the lineage. $n_{j\cdot\text{adj}}(\eta)$ denotes the number of *unique* values left in the p th attribute after adjusting for parents' splits.

Transitioning from the new tree back to the original tree involves pruning that node:

$$\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right) = \mathbb{P}(\text{PRUNE}) \mathbb{P}(\text{selecting } \eta \text{ to prune from}) = \mathbb{P}(\text{PRUNE}) \frac{1}{w_2^*}$$

where w_2^* denotes the number of second generation internal nodes (nodes with two terminal daughter nodes) in the new tree. Thus, the full transition ratio is:

$$\frac{\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right)}{\mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right)} = \frac{\mathbb{P}(\text{PRUNE})}{\mathbb{P}(\text{GROW})} \frac{b p_{\text{adj}}(\eta) n_{j\cdot\text{adj}}(\eta)}{w_2^*}.$$

Note that when there are no variables with more two or more unique values, the probability of GROW is set to zero and the step will be automatically rejected.

Likelihood Ratio

To calculate the likelihood, the tree structure determines which responses fall into which of the b terminal nodes. Thus,

$$\mathbb{P}\left(R_1, \dots, R_n \mid \mathfrak{T}, \sigma^2\right) = \prod_{\ell=1}^b \mathbb{P}\left(R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \sigma^2\right)$$

where each term on the right hand side is the probability of responses in one of the b terminal nodes, which are independent by assumption. The R_ℓ 's denote the data in the ℓ th terminal node and where n_ℓ denotes how many observations are in each terminal node and $n = \sum_{\ell=1}^b n_\ell$.

We now find an analytic expression for the node likelihood term. Remember, if the mean in each terminal node, which we denote μ_ℓ , was known, then we would have

$R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \mu_\ell, \sigma^2 \stackrel{iid}{\sim} \mathcal{N}(\mu_\ell, \sigma^2)$. BART requires μ_ℓ to be margined out, allowing the Gibbs sampler in Equation 3 to avoid dealing with jumping between continuous spaces of varying dimensions (Chipman et al., 2010, page 275). Recall that one of the BART model assumptions is a prior on the average value of $\mu \sim \mathcal{N}(0, \sigma_\mu^2)$ and thus,

$$\mathbb{P}\left(R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \sigma^2\right) = \int_{\mathbb{R}} \mathbb{P}\left(R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \mu_\ell, \sigma^2\right) \mathbb{P}\left(\mu_\ell; \sigma_\mu^2\right) d\mu_\ell$$

which can be shown via completion of the square or convolution to be

$$\begin{aligned} \mathbb{P}\left(R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \sigma^2\right) &= \frac{1}{(2\pi\sigma^2)^{n_\ell/2}} \sqrt{\frac{\sigma^2}{\sigma^2 + n_\ell\sigma_\mu^2}} \times \\ &\exp\left(-\frac{1}{2\sigma^2} \left(\sum_{i=1}^{n_\ell} (R_{\ell_i} - \bar{R}_\ell)^2 - \frac{\bar{R}_\ell^2 n_\ell^2}{n_\ell + \frac{\sigma^2}{\sigma_\mu^2}} + n_\ell \bar{R}_\ell^2\right)\right) \end{aligned} \quad (9)$$

where \bar{R}_ℓ denotes the mean response in the node and R_{ℓ_i} denotes the observations $i = 1 \dots n_\ell$ in the node.

Since the likelihoods are solely determined by the terminal nodes, the proposal tree differs from the original tree by only the selected node to be grown, denoted by ℓ , which becomes two daughters after the GROW step denoted by ℓ_L and ℓ_R . Hence, the likelihood ratio becomes:

$$\frac{\mathbb{P}(\mathbf{R} \mid \mathfrak{F}_*, \sigma^2)}{\mathbb{P}(\mathbf{R} \mid \mathfrak{F}, \sigma^2)} = \frac{\mathbb{P}(R_{\ell_{L,1}}, \dots, R_{\ell_{L,n_{\ell,L}}} \mid \sigma^2) \mathbb{P}(R_{\ell_{R,1}}, \dots, R_{\ell_{R,n_{\ell,R}}} \mid \sigma^2)}{\mathbb{P}(R_{\ell_1}, \dots, R_{\ell_{n_\ell}} \mid \sigma^2)} \quad (10)$$

Plugging Equation 9 into Equation 10 three times yields the ratio for the GROW step:

$$\sqrt{\frac{\sigma^2 (\sigma^2 + n_\ell \sigma_\mu^2)}{(\sigma^2 + n_{\ell_L} \sigma_\mu^2) (\sigma^2 + n_{\ell_R} \sigma_\mu^2)}} \exp \left(\frac{\sigma_\mu^2}{2\sigma^2} \left(\frac{(\sum_{i=1}^{n_{\ell_L}} R_{\ell_{L,i}})^2}{\sigma^2 + n_{\ell_L} \sigma_\mu^2} + \frac{(\sum_{i=1}^{n_{\ell_R}} R_{\ell_{R,i}})^2}{\sigma^2 + n_{\ell_R} \sigma_\mu^2} - \frac{(\sum_{i=1}^{n_\ell} R_{\ell,i})^2}{\sigma^2 + n_\ell \sigma_\mu^2} \right) \right)$$

where n_{ℓ_L} and n_{ℓ_R} denote the number of data points in the newly grown left and right daughter nodes.

Tree Structure Ratio

In Section 2.1 we discussed the prior on the tree structure (where the splits occur) as well as the tree rules. For the entire tree,

$$\mathbb{P}(\mathfrak{F}) = \prod_{\eta \in H_{\text{terminals}}} (1 - \mathbb{P}_{\text{SPLIT}}(\eta)) \prod_{\eta \in H_{\text{internals}}} \mathbb{P}_{\text{SPLIT}}(\eta) \prod_{\eta \in H_{\text{internals}}} \mathbb{P}_{\text{RULE}}(\eta)$$

where $H_{\text{terminals}}$ denotes the set of terminal nodes and $H_{\text{internals}}$ denotes the internal nodes.

Recall that the probability of splitting on a given node η is $\mathbb{P}_{\text{SPLIT}}(\eta) = \alpha / (1 + d_\eta)^\beta$. The probability is controlled by two hyperparameters, α and β , and d_η is the depth (number of parent generations) of node η . When assigning a rule, recall that **BART** picks from all available attributes and then from all available unique split points. Using the notation from the transition ratio section, $\mathbb{P}_{\text{RULE}}(\eta) = 1/p_{\text{adj}}(\eta) \times 1/n_{j\text{-adj}}(\eta)$.

Once again, the original tree features a node η that was selected to be grown. The proposal tree differs with two daughter nodes denoted η_L and η_R . We can now form the ratio:

$$\begin{aligned} \frac{\mathbb{P}(\mathfrak{F}_*)}{\mathbb{P}(\mathfrak{F})} &= \frac{(1 - \mathbb{P}_{\text{SPLIT}}(\eta_L)) (1 - \mathbb{P}_{\text{SPLIT}}(\eta_R)) \mathbb{P}_{\text{SPLIT}}(\eta) \mathbb{P}_{\text{RULE}}(\eta)}{(1 - \mathbb{P}_{\text{SPLIT}}(\eta))} \\ &= \frac{\left(1 - \frac{\alpha}{(1 + d_{\eta_L})^\beta}\right) \left(1 - \frac{\alpha}{(1 + d_{\eta_R})^\beta}\right) \frac{\alpha}{(1 + d_\eta)^\beta} \frac{1}{p_{\text{adj}}(\eta)} \frac{1}{n_{j\text{-adj}}(\eta)}}{1 - \frac{\alpha}{(1 + d_\eta)^\beta}} \end{aligned}$$

$$= \alpha \frac{\left(1 - \frac{\alpha}{(2+d_\eta)^\beta}\right)^2}{\left((1+d_\eta)^\beta - \alpha\right) p_{\text{adj}}(\eta) n_{j \cdot \text{adj}}(\eta)}$$

The last line follows from algebra and using the fact that the depth of the grown nodes is the depth of the parent node incremented by one ($d_{\eta_L} = d_{\eta_R} = d_\eta + 1$).

A.2 Prune Proposal

A prune proposal is the “opposite” of a grow proposal. Prune selects a node with two daughters and removes them. Thus, each ratio will be approximately the inverse of the ratios found in the previous section concerning the grow proposal. Note also that prune steps are not considered in trees that consist of a single root node.

Transition Ratio

We begin with transitioning from the original tree to the proposal tree:

$$\mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right) = \mathbb{P}(\text{PRUNE}) \mathbb{P}(\text{selecting } \eta \text{ to prune from}) = \mathbb{P}(\text{PRUNE}) \frac{1}{w_2}$$

where w_2 denotes the number of parent nodes that have two daughters but no granddaughters. To transition in the opposite direction, we are obligated to grow from node η . This is similar to Equation 8 except the proposed tree has one less terminal node due to the pruning of the original tree, resulting in a $1/(b-1)$ term:

$$\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right) = \mathbb{P}(\text{GROW}) \frac{1}{b-1} \frac{1}{p_{\text{adj}}(\eta^*)} \frac{1}{n_{j^* \cdot \text{adj}}(\eta^*)}.$$

Thus, the transition ratio is:

$$\frac{\mathbb{P}\left(\mathfrak{T}_* \rightarrow \mathfrak{T}\right)}{\mathbb{P}\left(\mathfrak{T} \rightarrow \mathfrak{T}_*\right)} = \frac{\mathbb{P}(\text{GROW})}{\mathbb{P}(\text{PRUNE})} \frac{w_2}{(b-1) p_{\text{adj}}(\eta^*) n_{j^* \cdot \text{adj}}(\eta^*)}.$$

Likelihood Ratio

This is simply the inverse of the likelihood ratio for the grow proposal:

$$\frac{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}_*, \sigma^2\right)}{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}, \sigma^2\right)} = \sqrt{\frac{(\sigma^2 + n_{\ell_L} \sigma_\mu^2) (\sigma^2 + n_{\ell_R} \sigma_\mu^2)}{\sigma^2 (\sigma^2 + n_\ell \sigma_\mu^2)}} \times \exp\left(\frac{\sigma_\mu^2}{2\sigma^2} \left(\frac{(\sum_{i=1}^{n_\ell} R_{\ell,i})^2}{\sigma^2 + n_\ell \sigma_\mu^2} - \frac{(\sum_{i=1}^{n_{\ell_L}} R_{\ell_L,i})^2}{\sigma^2 + n_{\ell_L} \sigma_\mu^2} - \frac{(\sum_{i=1}^{n_{\ell_R}} R_{\ell_R,i})^2}{\sigma^2 + n_{\ell_R} \sigma_\mu^2}\right)\right).$$

Tree Structure Ratio

This is also simply the inverse of the tree structure ratio for the grow proposal:

$$\frac{\mathbb{P}\left(\mathfrak{T}_*^{\mathfrak{A}}\right)}{\mathbb{P}\left(\mathfrak{T}^{\mathfrak{A}}\right)} = \frac{\left((1+d_\eta)^\beta - \alpha\right) p_{\text{adj}}(\eta^*) n_{j^* \cdot \text{adj}}(\eta^*)}{\alpha \left(1 - \frac{\alpha}{(2+d_\eta)^\beta}\right)^2}.$$

A.3 Change

A change proposal involves picking an internal node and changing its rule by picking both a new available predictor to split on and a new valid split value among values of the selected predictor. Although this could be implemented for use in any internal node in the tree, for simplicity we limit our implementation to *singly* internal nodes: those that have two terminal daughter nodes and thus, no grand-daughters.

Transition Ratio

The transition to a proposal tree is below:

$$\begin{aligned} \mathbb{P}\left(\mathfrak{T}^{\mathfrak{A}} \rightarrow \mathfrak{T}_*^{\mathfrak{A}}\right) &= \mathbb{P}(\text{CHANGE}) \mathbb{P}(\text{selecting node } \eta \text{ to change}) \times \\ &\quad \mathbb{P}(\text{selecting the new attribute to split on}) \times \\ &\quad \mathbb{P}(\text{selecting the new value to split on}) \end{aligned}$$

When calculating the ratio, the first three terms are shared in both numerator and denominator. The probability of selecting the new value to split on will differ as different split features have different numbers of unique values available. Thus we are left with

$$\frac{\mathbb{P}\left(\mathfrak{T}_*^{\mathfrak{A}} \rightarrow \mathfrak{T}^{\mathfrak{A}}\right)}{\mathbb{P}\left(\mathfrak{T}^{\mathfrak{A}} \rightarrow \mathfrak{T}_*^{\mathfrak{A}}\right)} = \frac{n_{j^* \cdot \text{adj}}(\eta^*)}{n_{j \cdot \text{adj}}(\eta)}$$

where $n_{j^* \cdot \text{adj}}(\eta^*)$ is the number of split values available under the proposal tree's splitting rule and $n_{j \cdot \text{adj}}(\eta)$ is the number of split values available under the original tree's splitting rule.

Likelihood Ratio

The proposal tree differs from the original tree only in the two daughter nodes of the selected change node. These two terminal nodes have the unexplained responses apportioned differently. Denote R_1 as the residuals of the first daughter node and R_2 as the unexplained responses in the second daughter node. Thus we begin with:

$$\frac{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}_*^{\mathfrak{A}}, \sigma^2\right)}{\mathbb{P}\left(\mathbf{R} \mid \mathfrak{T}^{\mathfrak{A}}, \sigma^2\right)} = \frac{\mathbb{P}\left(R_{1^*,1}, \dots, R_{1^*,n_{1^*}} \mid \sigma^2\right) \mathbb{P}\left(R_{2^*,1}, \dots, R_{2^*,n_{2^*}} \mid \sigma^2\right)}{\mathbb{P}\left(R_{1,1}, \dots, R_{1,n_1} \mid \sigma^2\right) \mathbb{P}\left(R_{2,1}, \dots, R_{2,n_2} \mid \sigma^2\right)}$$

where the responses denoted with an asterisk are the responses in the proposal tree's daughter nodes.

Substituting Equation 9 four times and using algebra, the following expression is obtained for the ratio:

$$\sqrt{\frac{\left(\frac{\sigma^2}{\sigma_\mu^2} + n_1\right) \left(\frac{\sigma^2}{\sigma_\mu^2} + n_2\right)}{\left(\frac{\sigma^2}{\sigma_\mu^2} + n_1^*\right) \left(\frac{\sigma^2}{\sigma_\mu^2} + n_2^*\right)}} \times \exp\left(\frac{1}{2\sigma^2} \left(\frac{\left(\sum_{i=1}^{n_1^*} R_{1^*,i}\right)^2}{n_1^* + \frac{\sigma^2}{\sigma_\mu^2}} + \frac{\left(\sum_{i=1}^{n_2^*} R_{2^*,i}\right)^2}{n_2^* + \frac{\sigma^2}{\sigma_\mu^2}} - \frac{\left(\sum_{i=1}^{n_1} R_{1,i}\right)^2}{n_1 + \frac{\sigma^2}{\sigma_\mu^2}} - \frac{\left(\sum_{i=1}^{n_2} R_{2,i}\right)^2}{n_2 + \frac{\sigma^2}{\sigma_\mu^2}}\right)\right)$$

which simplifies to

$$\exp\left(\frac{1}{2\sigma^2} \left(\frac{\left(\sum_{i=1}^{n_1^*} R_{1^*,i}\right)^2 - \left(\sum_{i=1}^{n_1^*} R_{1,i}\right)^2}{n_1 + \frac{\sigma^2}{\sigma_\mu^2}} + \frac{\left(\sum_{i=1}^{n_2^*} R_{2^*,i}\right)^2 - \left(\sum_{i=1}^{n_2^*} R_{2,i}\right)^2}{n_2 + \frac{\sigma^2}{\sigma_\mu^2}}\right)\right)$$

if the number of responses in the children do not change in the proposal ($n_1 = n_1^*$ and $n_2 = n_2^*$).

Tree Structure Ratio

The proposal tree has the same structure as the original tree. Thus we only need to take into account the changed node's daughters:

$$\frac{\mathbb{P}\left(\frac{\mathbb{P}}{\mathbb{P}}\right)_*}{\mathbb{P}\left(\frac{\mathbb{P}}{\mathbb{P}}\right)} = \frac{(1 - \mathbb{P}_{\text{SPLIT}}(\eta_{1^*})) (1 - \mathbb{P}_{\text{SPLIT}}(\eta_{2^*})) \mathbb{P}_{\text{SPLIT}}(\eta_*) \mathbb{P}_{\text{RULE}}(\eta_*)}{(1 - \mathbb{P}_{\text{SPLIT}}(\eta_1) (1 - \mathbb{P}_{\text{SPLIT}}(\eta_2))) \mathbb{P}_{\text{SPLIT}}(\eta) \mathbb{P}_{\text{RULE}}(\eta)}$$

The probability of splits remain the same because the daughter nodes are at the same depths. Thus we only need to consider the ratio of the probability of the rules. Once again, the probability of selecting the new value to split on will differ as different split features have different numbers of unique values available. We are left with

$$\frac{\mathbb{P}\left(\frac{\mathbb{P}}{\mathbb{P}}\right)_*}{\mathbb{P}\left(\frac{\mathbb{P}}{\mathbb{P}}\right)} = \frac{n_{j \cdot \text{adj}}(\eta)}{n_{j^* \cdot \text{adj}}(\eta^*)}$$

Note that this is the inverse of the transition ratio. Hence, for the change step, only the likelihood ratio needs to be computed to determine the Metropolis-Hastings ratio r .

B Bakeoff

We baked off nine regression data sets and assessed out-of-sample RMSE using 10-fold cross-validation. The results are displayed in Table 1.

We conclude that the implementation outlined in this paper performs approximately the same as the previous implementation with regards to predictive accuracy.

Dataset Name	bartMachine	BayesTree	randomForest
boston	4.66	4.59	4.58
triazine	0.04*	0.05	0.04
ozone	48.02	47.75	48.77
baseball	0.02	0.02*	0.02
wine.red	0.73	0.69*	0.76
ankara	2.62	2.61	3.46
wine.white	0.49	0.49	0.54
pole	4.00	3.65	11.56
compactiv	46.17*	47.92	40.74

Table 1: Average of 20 replicates of out-of-sample RMSE values on 9 datasets using three machine learning algorithms. Asterisks indicate a significant difference between **bartMachine** and **BayesTree** at a significance level of 5% with a Bonferroni correction. Comparisons with **randomForest**'s performance were not conducted.