# Bayesian Additive Regression Networks

Danielle Van Boxel[a,b]

[a]*Applied Math GIDP, University of Arizona, 617 N. Santa Rita, Room 410, Tucson, 85721, AZ, USA*
[b]*Data Diversity Lab, University of Arizona, Tucson, AZ, USA*

**Abstract**

We apply Bayesian Additive Regression Tree (BART) principles to training an ensemble of small neural networks for regression tasks. Using Markov Chain Monte Carlo, we sample from the posterior distribution of neural networks that have a single hidden layer. To create an ensemble of these, we apply Gibbs' sampling to update each network against the residual target value (i.e. subtracting the effect of the other networks). We demonstrate the effectiveness of this technique on several benchmark regression problems, comparing it to equivalent shallow neural networks, BART, and ordinary least squares. Our Bayesian Additive Regression Networks (BARN) provide more consistent and often more accurate results. On test data benchmarks, BARN averaged between 5% to 20% lower root mean square error. This error performance does come at the cost, however, of greater computation time. BARN sometimes takes on the order of a minute where competing methods take a second or less. But, BARN *without* cross-validated hyperparameter tuning takes about the same amount of computation time as tuned other methods. Yet BARN is still typically more accurate.

*Keywords:* BART, bayesian, ensemble, machine learning, MCMC

## 1. Introduction

In recent years, methods like Bayesian Additive Regression Trees (BART) have shown themselves capable of producing low prediction error on benchmark machine learning problems. BART uses an ensemble of decision trees, where each one is iteratively updated by sampling from a posterior distribution [4]. BART originally proved superior to modest neural networks in its time. Today, however, Neural Networks have come to the forefront of machine learning research [24]. Thus, we consider adapting the MCMC optimization tactics of BART into Bayesian Additive Regression Networks (BARN) in order to model even more accurately.

One powerful generalization of BART adapted it to a much broader set of problem types, not just regression and binary classification [15]. Standard BART relies on the particular structure of certain loss functions having closed-form posterior integrals in order to avoid the dimension changing from tree to tree. In a more recent study applying BART generically to problems like classification and count prediction, however, no closed form is required. The dimension changes as expected, but they account for this with Reversible Jump Markov Chain Monte Carlo (RJMCMC) [10]. With decision trees, it's possible to account for this in a general way (aided by considering only data points that bin into a particular leaf node). This expands the BART approach to new classes of machine learning problems.

One limitation, however, of using decision trees as the ensembled component in BART is that its output is inherently stepwise. That is, small changes in the input vector can lead to discontinuous changes in output. Even with an ensemble of trees and sampling over the posterior, the prediction function is not, in general, continuous. SoftBART addresses this by treating tree splits as random rather than deterministic

---

[16]. They note that this complicates the fitting procedure as SoftBART computes likelihoods over all leaf nodes rather than a single one. Another alternative to this is Model Trees BART (MOTR-BART) [20]. Here, one replaces the constant leaf output values with a small linear function of the inputs. MOTR-BART still has stepwise output, but it is stepwise linear rather than stepwise constant. And as this output is linear, the posterior still has a closed form via a Gaussian integral. MOTR-BART also helps reduce the number of trees. Each MOTR-BART tree can capture linear behavior, not just constant behavior. But in either model, the output can be continuous, matching many real problems we would like to model.

While neural networks dominate popular thinking in machine learning, especially with user interfaces like ChatGPT (and its predecessor, InstructGPT) [19], they have proven themselves in a variety of research contexts as well. In computer vision, for example, convolutional neural networks can provide not only high accuracy classification, but also a degree of explainability of portions of images that represent a particular class via techniques like Grad-CAM [25]. Researchers also make neural networks work together in ensembles. An interesting example of dynamic weighting of multiple networks helped address weather forecasting [17]. Such ensembling continues to see applied use more recently as well [18, 1, 26]. From all this, we can see that neural networks are a competitive algorithm in the ecosystem of machine learning approaches.

Decision trees as a backbone for BART are attractive because of the statistical rigor of a fully specified closed-form posterior. But with the increased demonstration of machine learning capabilities of neural networks [24], we go beyond MOTR-BART's extension of linear outputs [20] and replace decision trees entirely with neural networks. Using neural networks in this way changes the posterior calculation to no longer have a closed form. But we can still approximate the posterior and empirically evaluate such Bayesian Additive Regression Networks.

To adapt BART this way, we replace the standard posterior analysis with a revised one. We refer readers to Linero [15] for a detailed understanding of the MCMC fitting process of BART in a generic model-agnostic way. This provides background for Section 2, wherein we detail how BARN uses a MCMC fitting procedure inspired by BART. To show this empirically, Section 3 compares BARN against BART and other methods on some benchmark data sets. From this, we will see that BARN is highly adaptable and often more accurate than other regression methods.

## 2. BARN Model

We adapt the BART procedure [4], replacing the ensemble of decision trees with an ensemble of neural networks. Each neural network has a single hidden layer. That is, we train the ensemble such that $\hat{y}_i = \sum_k M_k(x_i)$, where $M_k$ is the $k$th neural network with some number of neurons and one hidden layer. In principle the model can be arbitrarily more complicated, but we retain this structure to simplify MCMC calculations, reduce computation time, and limit model overfitting. Due to the similarities with BART, we call our procedure Bayesian Additive Regression Networks (BARN).

At a high level, we propose architectural changes to a single model, $M_k$, at a time by computing the current residual, $R_k = Y - \sum_{j \neq k} M_j(X)$. This single model change enables us to condition on the fixed models, $M_{j \neq k}$, to effect Gibbs sampling. Each MCMC step proposes a modification to $M_k$'s architecture, trains this new model on $(R_k, X)$, and computes an acceptance probability, $A$. Note that because we determine the number of neurons within each network via the MCMC acceptance/rejection, this process implicitly performs a neural architecture search. Additionally, we fix the *number of models* in the ensemble as $N$. The MCMC procedure only modifies existing networks in the ensemble.

To do such modifications with detailed balance, we need the three components of the acceptance probability, $A$. For the transition probability, $T(M', M)$, we have some flexibility. While a "tilted" distribution would provide faster convergence, it also relies on more knowledge of the target distribution. As we generally lack such knowledge, we adopt a simple transition rule based on the number of neurons $m'$ and $m$ in the hidden layer of models $M'$ and $M$. Equation (1) specifies the probability of shrinking or growing a network by exactly 1 neuron. In our analysis, we set $p = 0.4$ as the growth transition probability. This helps regularize the networks to be small, helping avoiding overfitting. Finally, note that we do not allow networks to have $m = 0$ neurons; such transitions are always rejected.

$$T(M', M) = \begin{cases} p & m' = m + 1 \\ 1 - p & m' = max(1, m - 1) \end{cases} \tag{1}$$

Note that we do not consider the weight values themselves in this calculation; doing so would require handling dimensionality changes (i.e. with reversible jump MCMC [10]). And it would require some judgment or heuristic a priori about how the weights should change, *independent of the training process.* Also, with BART, the distribution of terminal node weights directly leads to a distribution on whole ensemble predictions (as the ensemble is a sum of these known terminal node distributions) [4]. In a neural network, however, computing the distribution of the output is more involved. Even if there is a distribution on weights, the distribution of input values is also required. And interaction between neurons is not independent, as they will have been trained with a traditional optimization. But such optimization is well studied and experimentally proven to be reliable [13] for neural networks, so we maintain such an approach for BARN. The rest of the procedure, however, mimics the BART MCMC approach.

When transitioning from $M$ to $M'$, we can make use of the already learned weights of $M$. If $m' = m + 1$, then we transfer the old weights to the new model and randomly initialize the new neuron weights. This retains previously learned information so the new model need not train exactly from scratch. Similarly, if $m' = m - 1$, then we retain the weights of the first $m - 1$ neurons. Note, however, that these serve only as the initial weights for the network.

Once the transition has been proposed and initial weights set, we can train this proposed network on the current training residual, $R_k$, and $X$. Note that one can use any number of neural network optimization methods; for simplicity we use the BFGS optimizer [7] with a maximum of 100 iterations (though this forms an additional hyperparameter to the BARN procedure).

Again, we reiterate that the development of a proposed replacement neural network, $M'$, is a two step process. First we mutate the architecture by adding or subtracting a node, $m' = m \pm 1$. And only then do we train the resulting architecture, informed by the original model weights, on the target residual, $R_k$. This produces a new set of weights, $w'_k$, associated with $M'$, and a new set of model predictions, $M'(X, w_k)$.

What remains is the posterior, comprised of the evidence and the prior. Recall that in BART, this is an integral over all possible output values for the given leaf. The BARN posterior breaks up in a similar way. The model node count prior, $P(M)$, does not depend on the weights, so it comes out of the integral. What remains within the integral are the error likelihood (i.e. evidence for the model) and the weight priors, seen simplified in Equation (2) and expanded with the normal error as well as flat weight assumptions in Equation (3).

$$P(R_k|M) = \int P(R_k|M_k, w)P(w)dw \tag{2}$$

$$P(R_k|M) = \int \frac{exp(\frac{(R_k - M_k(X, w))^2}{-2\sigma^2})}{\sqrt{2\pi\sigma^2}} 1 dw \tag{3}$$

Now we come to an issue. In BART or MOTR-BART, the analogous component is constant, or at least linear, in both cases lending itself to a closed-form integral [4, 20]. But $M_i(X, w)$ is a complicated nonlinear function of various $w_k$ weights, so we lack such a form. If we do away with the integral and make the weight sampling part of the MCMC algorithm, then the dimension of each model varies, requiring RJMCMC. Generalized BART handles something similar with reversible jumps, but they take advantage of the decision tree structure to again find a general closed form even under dimension change. The weights of a neural network, however, come from an involved training process, in our case BFGS. This makes obtaining a Jacobian for RJMCMC infeasible. We cannot compute a BARN posterior this way.

Instead, consider an approximation to the integral over weights. In neural networks, not unlike decision trees, certain weights produce much higher probabilities. That is, most of the probability mass for a given weight concentrates around a particular value. Additionally, swapping neuron positions (e.g. $w_{2i} \rightarrow w_{2i+1}$ and vice versa) produces numerically equivalent neural networks. These equivalent NNs must therefore have

multiple strong peaks in the multivariable probability density curves. Therefore, we approximate the integral itself with the likelihood of the peak (i.e. the integrand at the given weight values). This does introduce a potential scaling issue, but because we take a ratio of likelihoods (for the old network and the proposed network), these scales should cancel out, leaving us the approximate probability ratio. So we replace the integral of Equation (2) with the approximate Equation (4).

$$\int P(R_k|X, M_k)dw \approx \prod_{j \in \text{valid}} \frac{e^{-\frac{1}{2}\left(\frac{(R_k - M_k(X_j))}{\sigma}\right)^2}}{\sigma\sqrt{2\pi}} \tag{4}$$

Note that to compute this error likelihood, we apply the models to held-out data (but not the test data), else models may quickly overfit (an open question is if a prior on the error itself can better control overfitting without a held-out data set). This component then goes into the calculation of the acceptance ratio, $A$.

Next, we still need to specify a prior, $P(M)$, based on the neural network architecture as well as its learned weights. But as the weights can vary considerably depending on data, we restrict our prior to the size of the network. While different domains may propose different priors, we again want to encourage relatively small models. To that end, we encode a model's prior with a Poisson distribution on the count of neurons in its hidden layer. We choose $\lambda = 1$ as the mean of a Poisson distribution after empirically testing a few options (note that we explore more detailed cross-validation for hyperparameter selection in Section 3.1). So a model, $M$, with $m$ neurons has prior probability:

$$P(M) = \frac{\lambda^m e^{-\lambda}}{m!} \tag{5}$$

Finally, we combine the transition, evidence, and prior components into the acceptance ratio calculation. Equation (6) shows this for proposing a transition to $M_k'$ from $M_k$ with a residual of $R_k$.

$$A(M_k, M_k') = min\left(1, \frac{T(M_k|M_k')P(R_k|X, M_k')P(M_k')}{T(M_k'|M_k)P(R_k|X, M_k)P(M_k)}\right) \tag{6}$$

Recall that the residual in Equation (4) are sufficient to replace the contribution from all of the frozen models in the ensemble. When we fix those models and subtract their effect to obtain the residual, we are implementing a form of Gibbs sampling. That is, we approximate $P(Y|X, M_1, M_2, \cdots)$ as in [4] (i.e. the entire distribution given all the models, not just the particular model). Additionally, we sample the model error, $\sigma$, exactly as is done in BART, using an inverse gamma distribution. By iterating through the models in this fashion, we develop better and better approximations to the posterior. After some burn-in period, the distribution of models reflects the posterior given the available architectures and the data. We can then sample one (as for point inference) or more sequences of MCMC ensembles (as for intervals) as the model output.

## 3. Method Evaluation

We first assess the performance of BARN on benchmark regression problems that have been previously used in studies examining the performance of newly-developed methods. Then, we examine how BARN and competing methods perform on a series of controlled synthetic data sets. This allows us to see how BARN's performance changes in response to specific parameters like increasing noise or number of irrelevant features.

### 3.1. Model Evaluation Setup

To test the effectiveness of BARN, we train ensembles against nine benchmark regression problems and compare its error to that of existing regression techniques. Seven of the data sets analyzed in this study are freely available from the UCI Machine Learning repository [6], one is from a recent study [23], and the last is a randomly generated linear regression problem with eight relevant and two irrelevant variables. We compare BARN against BART which previously performed very well on some of these data sets [2]. Table 1 shows the size of each data set in features and data points.

Table 1: Representative data sets for testing BARN

| Dataset | Features | Data Points |
|---|---|---|
| boston | 13 | 506 |
| concrete | 9 | 1030 |
| crimes | 101 | 1994 |
| diabetes | 10 | 442 |
| fires | 10 | 517 |
| isotope | 1 | 700 |
| mpg | 7 | 398 |
| random | 10 | 1000 |
| wisconsin | 32 | 194 |

While these data sets provide a diverse set of problems on which to test BARN, we also implement a variety of synthetic data sets for additional testing. We generate this data by starting with different functional relationships, such as a randomly initialized linear relationship $y_i = \beta x_i$. Another function is pairs of indices of discrete clusters in the input space. We also extend the clusters by training a random forest on the indices and using its prediction as output. Finally, we have three multivariate functions from Friedman [8]. $F1$ is a 5-dimensional combined polynomial/trigonometric function used to evaluate BART [4]. $F2$ is a function involving products and reciprocals. Finally, $F3$ is similar to $F2$, but it is put through an arctangent function. We summarize these in Equation (7) - Equation (9). This diversity of relationships provides linear, nonlinear, and discontinuous data sets for modeling.

$$F1(X) = 10\sin(\pi X_1 X_2) + 20(X_3 - 1/2)^2 + 10X_4 + 5X_5 \tag{7}$$

$$F2(X) = \sqrt{X_1^2 + (X_2 X_3 - 1/(X_2 X_4))^2} \tag{8}$$

$$F3(X) = arctan\left(\frac{X_2 X_3 - 1/(X_2 X_4)}{X_1}\right) \tag{9}$$

$$\tag{10}$$

To further diversify our synthetic data, we perturb the output by adding varying amounts of noise (measured by signal-to-noise ratio, or SNR), adding extra irrelevant features to the $X$ vector, or changing the number of overall features. This creates problems of different complexity and challenge on which to test our methods. For more detail on this synthetic testing, Table 2 is a summary.

Table 2: Synthetic Data Generation Parameters

| Factors | Options |
|---|---|
| Functional Relationship | Cluster, Forest, Friedman (1-3), Linear |
| SNR | 10, 1 |
| # Relevant Features | 100, 10 |
| % Irrelevant Features Added | 90%, 10% |

For each data set, before training with our various algorithms, we do minimal preprocessing. For better comparison across data sets, we rescale outputs to be have mean 0 and variance 1. For input variables, we perform a principal component analysis (PCA) with the number of dimensions equal to the number of features (i.e. decorrelating but retaining all information). As we know neural networks benefit from this input transformation [14], we allow this step for all models as would be done in practical applications.

Now we specify the parameters needed for BARN. We set the number of networks in the ensemble to a small number, 10, even though BART recommends 100-200 trees [4]. Using only 10 networks instead of 100 helps reduce computation time since we have to train each network at every step. We also briefly

tested using BARN with a single neural network as a form of architecture search. This generally resulted in slightly worse models (about 2% worse RMSE than ensembled BARN of equivalent size), so we retain the ensemble approach. And as noted in Section 2, our growth transition probability is $p = 0.4$, and the prior on model size is a Poisson distribution with $\lambda = 1$ neuron. For training the NNs, we use BFGS optimization with up to 100 iterations and include a small $L_2$ regularization penalty of 0.001 on the weight magnitudes. To initialize the networks, we set the initial neuron count to 1 for all of them (mimicking BART's starting point of single node trees [4]), and train them each independently on $Y/10$ (i.e. assume they all contribute equally to the result). We run BARN for potentially 200 iterations, always doing at least 20. This is similar to the Keras `EarlyStopping` callback [5]. Every five steps, we check the validation error for improvement. If it gets worse by more than the tolerance, 0.0001, we stop training early per standard machine learning practice [9]. This helps ensure that models have converged. We accept the ensemble at the end of this procedure as the final one for testing.

To confirm that these MCMC parameters are acceptable, we apply a simple batch means analysis test. In our case, we measure $\phi_t = \sqrt{\sum_i \frac{(y_i - M^t(x_i))^2}{n}}$, the RMSE of the ensemble at each iteration, $t$. As we use the validation data for the acceptance testing, we also use it for this residual. Even with our modest amount of training, these values tended to be less than 1% the size of the RMSE variance across ensembles in independent runs, so we have some confidence that this level of burn-in is sufficient. Additionally, Figure 1 shows a typical error progression indicating convergence in RMSE. This type of convergence is often checked for "early stopping" of an algorithm [21]. Note that the RMSE converging means the error of the ensembles has converged, not that the parameters such as weights have converged within the MCMC. Because the networks transition, we don't necessarily expect the weights themselves to converge to a value. Finally, we are most interested in error estimates across data sets, so we limit discussing details about confirming convergence to just the overall RMSE approach.
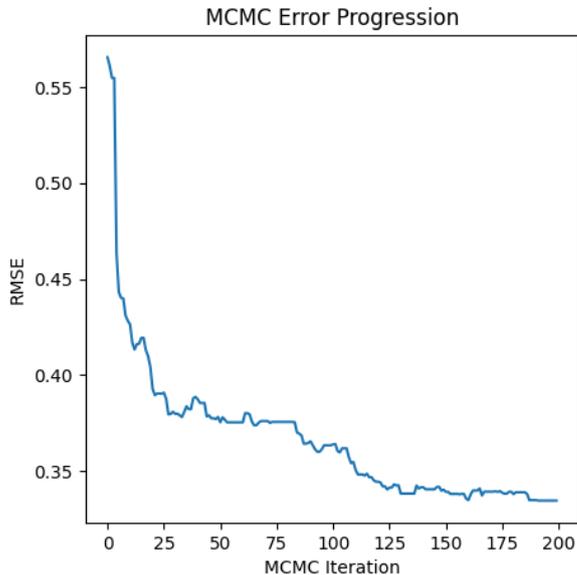


Figure 1: Typical error results for one of the data sets during the MCMC process shows burn-in achieved within the full run

For comparison to BARN, we train a single large neural network, a large BART forest, and an OLS model. For the large neural network, we retain a single hidden layer, but we set the number of neurons equal to the total neuron count in the BARN ensemble. Such a network has the same number of weights ((number of inputs+1) $\times$ number of neurons). We train this network for 2000 epochs using the same optimization and regularization parameters as used in BARN. For BART, we run with default parameters

with `wbart` in the BART package (100 trees, $k = 2, \nu = 3, q = 0.99$) as those have been shown to perform well [4]. This already encourages small trees, and BART is fast enough that the default burn-in of 1000 iterations is sufficient and computable. This provides three alternatives against which to evaluate BARN.

For each method (excluding OLS), we also test a version of it with five-fold cross-validated hyperparameter tuning (denoted with "CV" in tables and figures). For BARN, we explore choices for the prior mean of the model size, $\lambda$; the number of networks in the ensemble; and the activation function in the neural network. With two options for each, we have a total of eight configurations to test on each of five folds. For BART, we fix the number of trees at 100, but we vary $k$, $\nu$, and $q$ for a total of 12 combinations. With the single large neural network, we consider model sizes in different multiples of the BARN-derived size, different learning rates, number of epochs, and like BARN, two options for the activation function. Table 3 summarizes the options for each method. In all cases, we select the set of parameters with the smallest root mean square error.

Table 3: Hyperparameter options for cross-validated methods

| Method | Parameter | Values |
|--------|-----------|--------|
| BARN | $n$, Number of Networks | 10, 20 |
| | $\lambda$, Poisson Mean Number of Neurons Prior | 1, 2 |
| | NN Activation Function | Sigmoid, ReLU |
| BART | $k$, standard deviation prior on $\sigma_\mu$ | 1,2,3,4,5 |
| | $(\nu, q)$, priors on $\sigma$ | $((3,0.9), (3,0.99), (10,0.75))$ |
| Big NN | BARN neuron count multiplier | 1, 2, 10 |
| | $lr$, Adam initial learning rate | 1e-5, 1e-4 |
| | Number of epochs | 2000, 4000 |
| | NN Activation Function | Sigmoid, ReLU |

For each data set, we perform 40 random trials with each type of model. In each trial, we randomly split the data into 50% training, 25% validation (for early stopping checking), and 25% testing, using the same split across model types for consistency. After completion, we compute the average RMSE of the test predictions.

### 3.2. Benchmark Results

Table 4 shows the mean RMSE for each data set and model. We summarize this across data sets into the relative RMSE in Figure 2, and we compile the median points for both absolute and relative RMSE into Table 5, along with pooled sample standard deviations. Each point in Figure 2 is the test RMSE of that method divided by the test RMSE of the best method for that particular data split. Additionally, Table 6 shows the maximum relative error seen for each model and data set. Note that for BARN, the forest fires data set is the only one with relative error large enough to contain the more severe outliers in Figure 2. Finally, Figure 3 display $R^2$ values of methods averaged across data sets, with error bars showing the pooled variance estimate. Because different data sets have different amounts of noise, we expect them to have different *mean* $R^2$ values. By using pooled variance, that is, the variance about *each* mean, we can estimate the variance expected in a typical data set. Ignoring this and using sample variance would artificially inflate the error bars. From these metrics, we see that BARN appears to have generally lower relative error and better fit across data sets.

For the synthetic data sets, we perform a similar analysis. Here, we focus on how changing each parameter affects the results. We do, however, include the overall relative error and $R^2$ in Figure 4a, where BARN again outcompetes alternatives. For more detail, we have a set of subplots for each of the variants in Figure B.6 in Appendix Appendix B. Across all of these, we can further inspect BARN's robustness.

By examining a broad cross section of real and synthetic data, we can compare BARN to other methods in highly variable situations. Some of these, like the forest fires data set, have non-normal error terms. Others, like the Wisconsin breast cancer set, have a very limited amount of data available for training and testing. What's more, we see from the synthetic data how BARN responds to changing the amount of noise

Table 4: Mean scaled RMSE over 40 trial with smallest mean error in bold for each data set

| Data set | BARN | BARN CV | BART | BART CV | Big NN | Big NN CV | OLS |
|---|---|---|---|---|---|---|---|
| california | 0.513 | **0.495** | 0.524 | 0.498 | 0.857 | 0.566 | 0.558 |
| concrete | 0.375 | **0.367** | 0.533 | 0.462 | 6.266 | 0.702 | 0.631 |
| crimes | 0.622 | 0.627 | 0.708 | 0.649 | 0.736 | **0.599** | 0.614 |
| diabetes | 0.734 | 0.745 | 0.807 | 0.791 | 0.721 | 0.724 | **0.719** |
| fires | 1.494 | 1.441 | 1.358 | 1.221 | 4.088 | 1.130 | **1.113** |
| isotope | 0.299 | **0.298** | 0.327 | 0.315 | 0.322 | 0.329 | 0.299 |
| mpg | **0.354** | 0.361 | 0.480 | 0.454 | 19.539 | 0.461 | 0.408 |
| random | 0.073 | 0.074 | 0.386 | 0.179 | 0.096 | 0.089 | **0.068** |
| wisconsin | 1.123 | 1.097 | 1.104 | 1.058 | 9.028 | **1.024** | 1.189 |
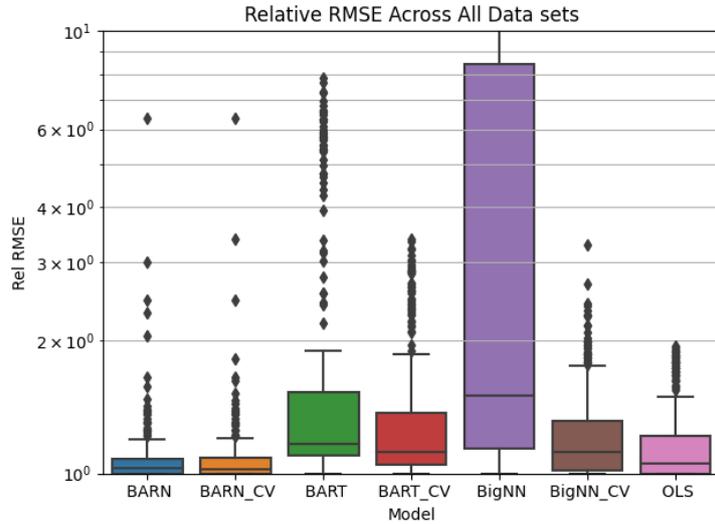


Figure 2: BARN is more adaptable to different problems than other methods and outperforms them across data sets

and presence of many features (irrelevant or not). As visible in the figures, BARN appears competitive with other methods, often outperforming or matching the next-best method.

### 3.3. Discussion

Based on these visualizations, we see that BARN tends to have lower testing error than shallow neural networks, BART, and OLS across data sets. We choose to recompute BART results rather than using those in Biau et al. [2] in order to maintain data set and data split consistency. The results in Table 4 show BARN is often the best performing, with either BARN or tuned BARN ("BARN_CV") having the lowest RMSE in four of nine data sets. In two of these, california and isotope, however, other methods like tuned BART or OLS were within 1% of BARN's error. Note that it is possible for the non-cross-validation version to have lower testing error, as the model selection is done on training data. The relative RMSE shown in Figure 2 displays even more clearly how BARN adapts well across a variety of problems, with BARN having tighter and lower errors than other methods. Table 5 makes this quantitative; on real data we see that BARN (even without tuning) has a median testing error about 13% lower than the next-best method, *which is tuned.* While BART or OLS may dominate a particular problem, BARN has the smallest relative error across data and a much tighter interquartile range. It does have a seemingly large number of outliers (similar to BART without tuning), but Table 6 indicates that these large errors come from the forest fires data set, as the maximum relative error for BARN on other data sets is smaller than the outliers. In the forest fire data set,

Table 5: Median test RMSE, relative RMSE, and pooled standard deviation across types of data sets, excluding forest fires data

| Data set, metric | BARN | BARN CV | BART | BART CV | Big NN | Big NN CV | OLS |
|---|---|---|---|---|---|---|---|
| Real, Median | 0.390 | **0.383** | 0.517 | 0.477 | 0.732 | 0.570 | 0.562 |
| Real, pooled $\sigma$ | 0.071 | 0.057 | 0.056 | 0.056 | 9.634 | 0.070 | 0.059 |
| Real, rel, Median | 1.016 | **1.014** | 1.168 | 1.117 | 1.409 | 1.170 | 1.089 |
| Real, rel, pooled $\sigma$ | 0.096 | 0.062 | 0.391 | 0.156 | 25.789 | 0.173 | 0.079 |
| Synth | 0.788 | **0.761** | 0.881 | 0.847 | 0.873 | 0.791 | 0.780 |
| Synth, pooled $\sigma$ | 0.083 | 0.082 | 0.068 | 0.073 | 0.078 | 0.118 | 0.059 |
| Synth, rel | 1.050 | **1.026** | 1.206 | 1.158 | 1.207 | 1.053 | 1.060 |
| Synth, rel, pooled $\sigma$ | 0.098 | 0.089 | 0.107 | 0.105 | 0.115 | 0.253 | 0.084 |

Table 6: Maximum Relative RMSE for each modeling method and data set

| Data set | BARN | BARN CV | BART | BART CV | Big NN | Big NN CV | OLS |
|---|---|---|---|---|---|---|---|
| california | 2.31 | 1.09 | 1.21 | 1.11 | 10.96 | 1.34 | 1.32 |
| concrete | 1.17 | 1.13 | 1.72 | 1.55 | 76.37 | 3.28 | 1.94 |
| crimes | 1.04 | 1.06 | 1.23 | 1.09 | 1.40 | 1.00 | 1.04 |
| diabetes | 1.19 | 1.44 | 1.29 | 1.25 | 1.10 | 1.15 | 1.07 |
| fires | 6.34 | 6.34 | 3.36 | 2.09 | 25.94 | 1.99 | 1.21 |
| isotope | 1.00 | 1.00 | 1.15 | 1.08 | 1.28 | 1.82 | 1.00 |
| mpg | 1.10 | 1.40 | 1.67 | 1.57 | 354.35 | 1.76 | 1.37 |
| random | 1.24 | 1.36 | 7.82 | 3.40 | 1.98 | 1.86 | 1.01 |
| wisconsin | 1.39 | 1.39 | 1.36 | 1.36 | 48.79 | 1.26 | 1.56 |

the target value is the area of forest burned in a given time period. Many periods have zero, while others vary, making it zero-inflated or zero-censored data. Other methods also face challenges with this data, but not always as severely. Finally, Figure 3 again shows how BARN finds better fits (in both training and test data) than other approaches. BARN, tuned or not, has about mean 0.6 $R^2$, while the next highest is tuned BART at 0.58. These $R^2$ values, however, are within the pooled standard deviation, close enough to not be significantly different. So although BARN is not universally superior, there is a broad spectrum of situations in which it noticeably outperforms existing methods and may be more generally robust.

We also call out some potential peculiarities in this analysis. First, the large single NN, without cross-validated tuning, performed very poorly on the real data, with median relative testing error over 50% higher than BARN. The available default parameters often failed to converge on various problems, leading to extremely poor results (e.g. "going off the edge of the plot" in Figure 2). While one could devote more effort to finding better defaults, from the cross-validated tuning, we see that it can perform reasonably well. And generally, one set of parameters may not transfer well between problem domains. For BARN, however, cross-validated tuning provides very little advantage. The view in Figure 3 perhaps makes this most clear. There is a small boost in *training* $R^2$ with BARN CV (about 0.07 higher than plain BARN), but this does not lead to much of a boost in *testing* $R^2$ (less than 0.01 improvement). This increase is perhaps just enough to edge out the default parameters as seen in Table 4, but that tuning may not worth the effort. Both BARN CV and regular BARN are have $R^2$ values well within a pooled standard deviation of each other. Cross-validation on BARN set the mean model size prior to be $\lambda = 1$, which was the default choice, about 85% of the time. Combining these observations about BARN and the large NN provides an additional insight. A BARN ensemble of shallow neural networks can be trivially rewritten as a single neural network that takes a weighted sum of all the neurons in all the networks. That is, they are structurally identical. Yet, BARN's training procedure provides it with robustness comparable to that of cross-validation for the the large NN. Large NNs have the opportunity to numerically duplicate BARN's results, but BARN is generally superior.

Using the synthetic data sets, we see similar behavior in BARN. As Figure 4a shows, BARN again has a small relative test RMSE across data sets. We observe two additional points here. First, cross-validated
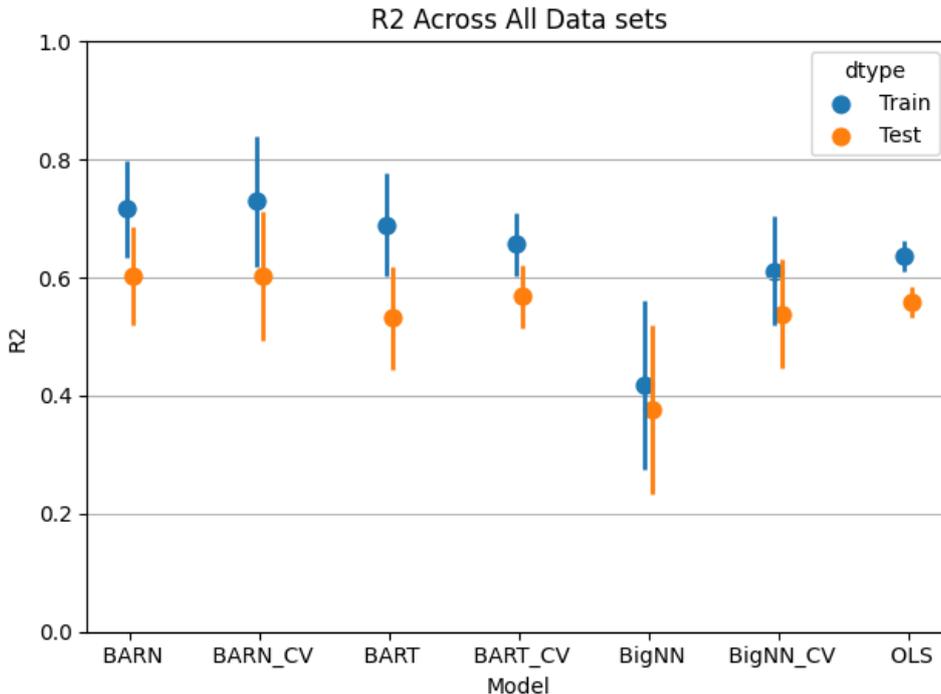
Figure 3: BARN resists overfitting similar to other methods when comparing $R^2$ values across training and testing data with pooled variance

hyperparameter tuning does have a modest impact on BARN (about 0.01 improvement on testing data), though perhaps still less so than on other methods. Even without, BARN is competitive with other methods that use such tuning. Its mean relative test RMSE was 1.1, and with tuning only 1.08. The next lowest mean relative test RMSE belonged to the tuned NN at 1.15, about 5% worse than either BARN method. Turning to Figure 4b, we can inspect how BARN handles overfitting on our synthetic data sets. It has a similar amount of overfitting compared with other nonlinear methods, with an $R^2$ about 0.3 lower on testing data than training. Not surprisingly, OLS overfits less, with a decrease in $R^2$ of about 0.2. But BARN, with our without tuning, has the highest mean $R^2$, though error bars overlap with tuned neural networks and OLS. Overall, BARN still performs well on these synthetic sets.

Examining performance on specific sets, we find a few interesting and distinct results. For the cluster and forest data sets (which mimic an underlying discontinuous relationship), BARN performs better than other methods when there are more features and noise. Tuned BARN has a mean RMSE of about 0.8 in these situations, whereas the competing methods are 0.9 or higher. The testing $R^2$ reflects this as well, where tuned BARN is about 0.35 (still poor, but higher than others). When $SNR = 10$ (i.e. low noise) for these data sets, however, other approaches like neural networks have lower error and higher $R^2$, indicating that perhaps BARN has robustness in this situation, but may not be the most accurate when the problem is easier. The Friedman data sets have somewhat different perspective. For $F2$ and $F3$, BARN seems better overall, regardless of noise, feature count, or irrelevant features. With $F1$, it is only superior for $SNR = 10$ cases, especially when there are fewer irrelevant features. Note that this is the opposite of what we see with the cluster and forest data. We speculate that the periodic nonlinearity of $F1$ is difficult for BARN to track under noisy conditions. Finally, we note that in the linear data sets, OLS generally performs best, as expected. More detailed subplots demonstrating these finds are available in Figure B.6. Having these synthetic data sets where we know the underlying parameters enables to see that BARN is robust under a variety of conditions where other approaches may have difficulty.

From all these plots and analyses, we see that BARN is effective compared to other methods on many different domains. We briefly note here, however, that this performance does take additional computation time. While we save detailed discussion of timing results for another study, BARN without cross-validated hyperparameter tuning takes comparable time to BART or neural networks *with* such tuning. Yet, even the untuned BARN shows low testing errors in many situations. It brings together the strengths of both BART-style MCMC model search and modern neural network implementation to produce a better machine learning method.

## 4. Summary

Though Section 3 shows BARN to be competitive, there remains a significant amount of future work on better theoretical understanding and practical implementation. Most of the theoretical advances relate to more rigorous analysis of the MCMC components. For example, our selection of transition probabilities was somewhat arbitrary. While this limitation was also in Bayesian CART [3] regarding tree growth and pruning, a follow-up study on generalized BART derived rigorous normal distributions for transition probabilities [15]. We may be able to derive a similar distribution to improve convergence and reduce variance in the produced models. Similarly, we need better justification for our model size prior. While a Poisson or negative binomial distribution is probably sufficient, a heuristic to guide both the distribution and parameters would be helpful. From cross-validation, we found a mean of $\lambda = 1$ for our Poisson prior was most common, so we recommend this as a starting point. Next, we assumed the errors for each model are independently identically distributed normal, but this is problem dependent. One can account for relaxations to this by modifying the $P(Y|X, M)$ expression, but it is not clear if we can derive an approximation for this similar to BART's closed forms. And finally, we should seek either an RJMCMC-compatible form of model proposal with changing dimensions or a better justification for avoiding this. The approximation of Equation (4) was a start, but one may be able to better account for changing likelihood functions. Alternatively, a fixed-dimension MCMC may become feasible by setting a maximum network size and keeping most of the weights zero (effectively deleting most of the neurons). Or, one might propose weights entirely via the MCMC process, completely eschewing traditional weight optimization approaches entirely (as BART does). We believe our theoretical approach is sound and confirmed by the experimental results, but this remains an open area of research.

On the applied side, there is also fruitful future research. Still somewhat theoretical, but with direct applications, is a version of BARN suitable for classification problems. BART provides an approach to this using binary probit [4], which we can adapt to BARN. Next, making BARN more accessible via simplified software interfaces and additional languages will help it reach more researchers and data scientists. This also means adapting BARN to a wider variety of backbones. To illustrate: BART uses decision trees and BARN uses neural networks, but a generic Bayesian Additive Regression *Model* could use, say, support vector machines (SVMs) [11]. SVMs, in particular, are attractive as they may have fewer dimension-changing issues if a single type of kernel is considered. An alternative different application is to focus more on the architecture search rather than the ensembling. If we restrict the number of networks to exactly one, then BARN transition proposals and MCMC acceptance sample from the posterior on the total number of neurons. With a single hidden layer in the neural net, such an approach may be less effective than exhaustively trying many different neuron counts. But if we expand the neural network architecture to allow proposing additional layers, or even different types of layers, such as convolutional ones in an image processing context, a BARN-guided search may become more competitive than exhaustion or even genetic algorithm methods [12]. Such extensions to BARN may make it more applicable on ever more diverse problems.

We have shown BARN to accurately model a variety of regression problem domains. By adapting the MCMC model sampling process of BART to modern neural networks, we employ the strengths of both. Though our posterior is an approximation, empirical benchmarking shows how BARN can achieve lower test errors on real data sets. We further demonstrate this effectiveness on synthetic data, setting various parameters like noise and number of features, which are not controllable in real data. BARN may take additional processing time, but it may not require much hyperparameter tuning, making its overall computation time comparable to other methods like BART and pure neural networks. Therefore, BARN is

a machine learning approach we should consider applying to new data science problems when seeking lower error rates.

## Code Availability

Software for running BARN in general is available under `barmpy` on PyPi [22]. Code specific to our analysis is available on request to the author.

## Acknowledgements

## Appendix A. Benchmark Regression Posterior

Here we briefly review the neuron count distributions within the BARN ensemble on different types of data. Figure A.5 shows the posterior distribution of neuron counts for both the tuned and untuned BARN methods. In addition, we include on each graph the prior distribution (with the caveat that zero neurons was never actually possible). For all but the concrete data set, the final distribution most often has a single neuron in each network (about 80% of the time). The concrete data is highly nonlinear, so it not surprising that BARN needs more neurons (2 to 4) per network to model this.

## Appendix B. Synthetic Regression Details

We include some additional figures which go into more detail for individual data sets. Figure B.6 shows the relative testing and training error results for each synthetic data set, organized by the factors that comprise them.
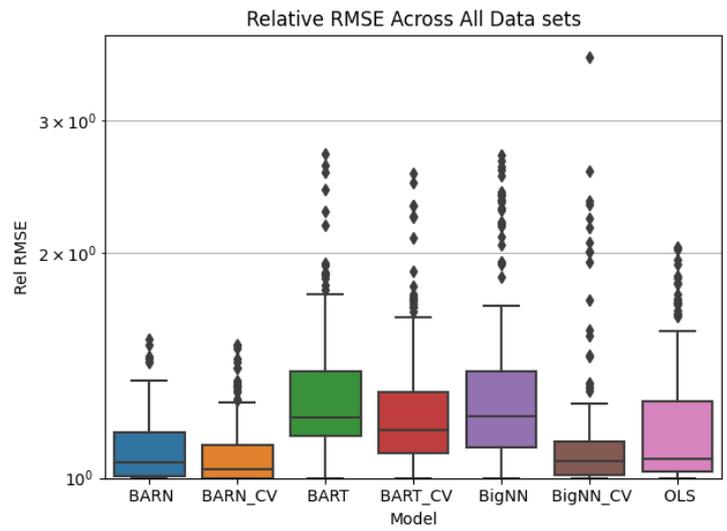
(a) BARN is adaptable, as seen on a variety of synthetic data sets
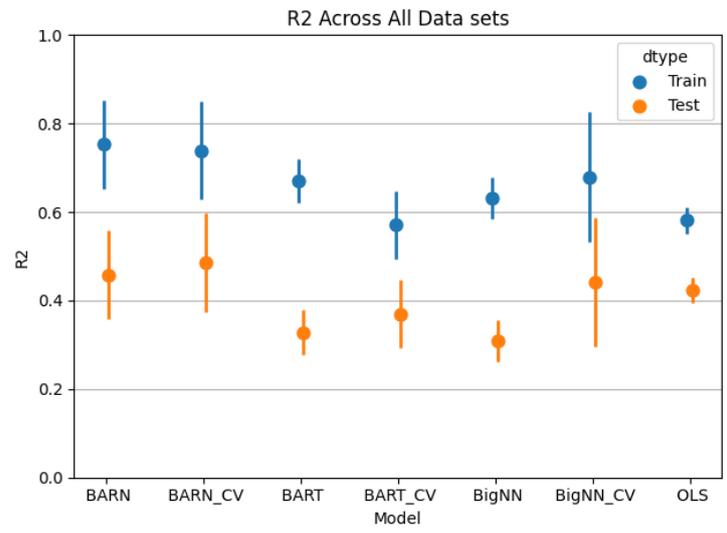


(b) $R^2$ for synthetic training and testing data shows some overfit

Figure 4: BARN is competitive with other methods across data sets

13
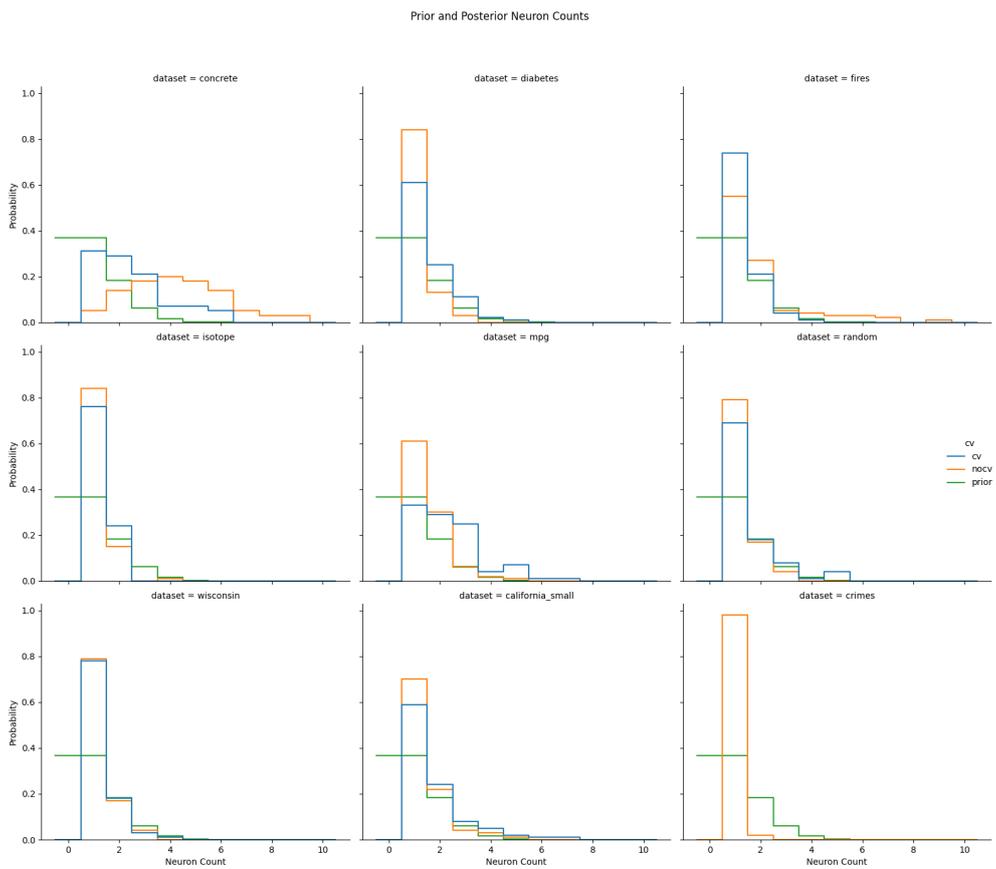
Figure A.5: Posterior distribution of neuron counts varies across data sets
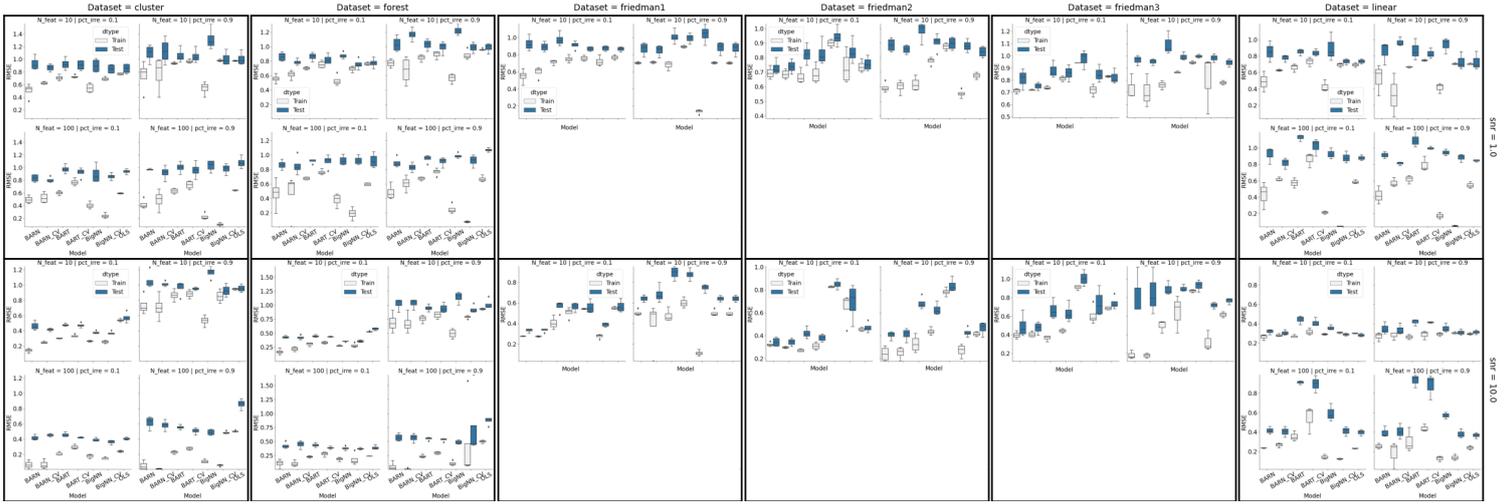
(a) Individual RMSE results, training and test
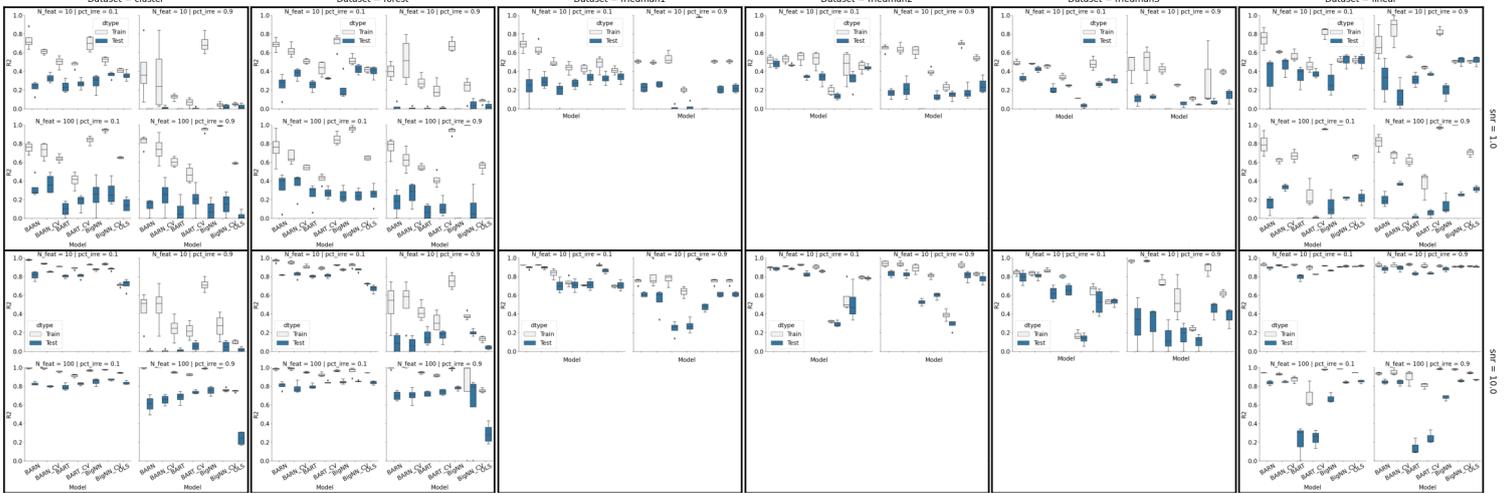


(b) Individual $R^2$ results, both training and test

Figure B.6: Synthetic data results, grouped by data set and SNR level (heavy boxes), as well as number of features and irrelevant features (interior division)

# References

[1] Alam, K.M.R., Siddique, N., Adeli, H., 2020. A dynamic ensemble learning algorithm for neural networks. Neural Computing and Applications 32, 8675–8690.

[2] Biau, G., Scornet, E., Welbl, J., 2019. Neural random forests. Sankhya A 81, 347–386.

[3] Chipman, H.A., George, E.I., McCulloch, R.E., 1998. Bayesian CART model search. Journal of the American Statistical Association 93, 935–948.

[4] Chipman, H.A., George, E.I., McCulloch, R.E., et al., 2010. BART: Bayesian additive regression trees. The Annals of Applied Statistics 4, 266–298.

[5] Chollet, F., et al., 2015. Keras. https://keras.io.

[6] Dua, D., Graff, C., 2017. UCI machine learning repository. URL: http://archive.ics.uci.edu/ml.

[7] Fletcher, R., 2000. Practical methods of optimization. John Wiley & Sons.

[8] Friedman, J.H., 1991. Multivariate adaptive regression splines. The annals of statistics 19, 1–67.

[9] Gençay, R., Qi, M., 2001. Pricing and hedging derivative securities with neural networks: Bayesian regularization, early stopping, and bagging. IEEE transactions on neural networks 12, 726–734.

[10] Green, P.J., 1995. Reversible jump markov chain monte carlo computation and bayesian model determination. Biometrika 82, 711–732.

[11] Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H., 2009. The Elements Of Statistical Learning: Data Mining, Inference, And Prediction. volume 2. Springer.

[12] Idrissi, M.A.J., Ramchoun, H., Ghanou, Y., Ettaouil, M., 2016. Genetic algorithm for neural network architecture optimization, in: 2016 3rd International Conference on Logistics Operations Management (GOL), IEEE. pp. 1–4.

[13] Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .

[14] LeCun, Y., Bottou, L., Orr, G., Müller, K., 1998. Efficient backprop, in: Orr, G.B., Müller, K.R. (Eds.), Neural networks: tricks of the trade, Springer.

[15] Linero, A.R., 2022. Generalized bayesian additive regression trees models: Beyond conditional conjugacy. arXiv preprint arXiv:2202.09924 .

[16] Linero, A.R., Yang, Y., 2018. Bayesian regression tree ensembles that adapt to smoothness and sparsity. Journal of the Royal Statistical Society Series B: Statistical Methodology 80, 1087–1110.

[17] Maqsood, I., Khan, M.R., Abraham, A., 2004. An ensemble of neural networks for weather forecasting. Neural Computing & Applications 13, 112–122.

[18] Mongan, P.G., Hinchy, E.P., O'Dowd, N.P., McCarthy, C.T., Diaz-Elsayed, N., 2023. An ensemble neural network for optimising a cnc milling process. Journal of Manufacturing Systems 71, 377–389.

[19] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al., 2022. Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems 35, 27730–27744.

[20] Prado, E.B., Moral, R.A., Parnell, A.C., 2021. Bayesian additive regression trees with model trees. Statistics and Computing 31, 1–13.

[21] Prechelt, L., 2002. Early stopping-but when?, in: Neural Networks: Tricks of the trade. Springer, pp. 55–69.

[22] PyPi Maintainers, 2023. Python package index - PyPi. https://pypi.org/.

[23] Román Palacios, C., Carroll, H., Arnold, A., Flores, R., Petersen, S., McKinnon, K., Tripati, A., Gan, Q., 2022. Bayclump: Bayesian calibration and temperature reconstructions for clumped isotope thermometry. Authorea Preprints .

[24] Schmidhuber, J., 2015. Deep learning in neural networks: An overview. Neural networks 61, 85–117.

[25] Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE international conference on computer vision, pp. 618–626.

[26] Vallabhajosyula, S., Sistla, V., Kolli, V.K.K., 2022. Transfer learning-based deep ensemble neural network for plant leaf disease detection. Journal of Plant Diseases and Protection 129, 545–558.