

RESEARCH

Open Access



A flexible approach for variable selection in large-scale healthcare database studies with missing covariate and outcome data

Jung-Yi Joyce Lin¹, Liangyuan Hu^{2*}, Chuyue Huang³, Ji Jiayi², Steven Lawrence¹ and Usha Govindarajulu¹

Abstract

Background: Prior work has shown that combining bootstrap imputation with tree-based machine learning variable selection methods can provide good performances achievable on fully observed data when covariate and outcome data are missing at random (MAR). This approach however is computationally expensive, especially on large-scale datasets.

Methods: We propose an inference-based method, called RR-BART, which leverages the likelihood-based Bayesian machine learning technique, Bayesian additive regression trees, and uses Rubin's rule to combine the estimates and variances of the variable importance measures on multiply imputed datasets for variable selection in the presence of MAR data. We conduct a representative simulation study to investigate the practical operating characteristics of RR-BART, and compare it with the bootstrap imputation based methods. We further demonstrate the methods via a case study of risk factors for 3-year incidence of metabolic syndrome among middle-aged women using data from the Study of Women's Health Across the Nation (SWAN).

Results: The simulation study suggests that even in complex conditions of nonlinearity and nonadditivity with a large percentage of missingness, RR-BART can reasonably recover both prediction and variable selection performances, achievable on the fully observed data. RR-BART provides the best performance that the bootstrap imputation based methods can achieve with the optimal selection threshold value. In addition, RR-BART demonstrates a substantially stronger ability of detecting discrete predictors. Furthermore, RR-BART offers substantial computational savings. When implemented on the SWAN data, RR-BART adds to the literature by selecting a set of predictors that had been less commonly identified as risk factors but had substantial biological justifications.

Conclusion: The proposed variable selection method for MAR data, RR-BART, offers both computational efficiency and good operating characteristics and is utilitarian in large-scale healthcare database studies.

Keywords: Missing at random, Multiply imputed datasets, Tree-based methods, Variable importance

*Correspondence: liangyuan.hu@rutgers.edu

²Department of Biostatistics and Epidemiology, Rutgers University, 683 Hoes Lane West, 08854 Piscataway, USA

Full list of author information is available at the end of the article



© The Author(s). 2022 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

Background

The problem of variable selection is one of the most popular model selection problems in statistical applications [1]. Variable selection involves modeling the relationships between an outcome variable and a set of potential explanatory variables or predictors and identifying a subset of predictors that has the most impact on the model fit. Variable selection in the presence of missing data has gained growing attention in recent years, particularly among statistical practitioners working with health datasets, which frequently present the missing data issue.

There are three general missing data mechanisms: missing completely at random (MCAR), missing at random (MAR), and missing not at random (MNAR) [2]. When the missingness depends neither on observed data nor on the missing data, the data are said to be MCAR. When data are MCAR, the complete cases are a random subsample of the population, thus results from a complete cases analysis would not be biased but can be less efficient when a large amount of data are discarded. A more realistic assumption about the missing data is that the mechanism of missingness may depend on the observed data, and then the missing data are MAR given the observed data [3]. Under the MAR assumption, one can impute the missing values based on the observed data. In a more challenging situation where the missingness depends on the missing data, the data are MNAR [4]. To handle MNAR, an approach recommended by the National Research Council [5] is sensitivity analysis [6, 7], which evaluates the impact of the potential magnitude of departure from MAR on analysis results. Little et al. [4] provide a comprehensive review of existing statistical approaches for handling missing data. In this article, we focus on the MAR mechanism that allows replacing missing data with substituted values or imputation based on the observed data, which is widely accepted in epidemiological and health research.

Under MAR, variable selection can be conducted in combination with imputation for missing values [8]. This approach is conceptually straightforward and is widely applicable to general missing data patterns [8]. The key consideration is how to combine the uncertainties about the imputation and the selection of predictors.

Wood et al. [9] compared strategies for combining the backward stepwise selection approach and multiple imputation for incomplete data, and recommended selecting variables based on Rubin's rules that combine estimates of parameters and standard errors across multiple imputed data. Long and Johnson [8] proposed combining bootstrap imputation and stability selection [10]. Bootstrap imputation is also known as "BS-then-MI". The term refers to bootstrapping first, and then applying the imputation method within each bootstrap sample. In Long and Johnson [8], single imputation is performed on each bootstrap sample, and then in stability selection, the ran-

domized lasso is applied to each of M bootstrap samples; a set of predictors are deemed as important if they are selected in at least πM samples, where π is a proportion representing the selection threshold. Both studies rely on the parametric assumptions, in which the exact relationships between response and covariates need to be made explicit.

Flexible nonparametric methods can mitigate the reliance on the parametric assumptions and improve the variable selection results [11–18]. A recent work by Hu et al. [19] investigated a general strategy that combines bootstrap imputation with six variable selection methods (four tree-based and two parametric) for variable selection among incomplete data. Their numerical studies suggest that flexible tree-based methods achieved substantially better variable selection results than parametric methods with MAR data in conditions of nonlinearity and nonadditivity. Among the four tree-based methods, permutation-based Bayesian additive regression trees (BART) [11] and recursive elimination based extreme gradient boosting (XGBoost) [20] were the top performers. We term these two methods as BI-BART and BI-XGB, against which we will benchmark our proposed method for variable selection in the presence of MAR data. Although BI-BART and BI-XGB have good performance, they are computationally expensive as they require both bootstrap and iterative variable selection procedures on each bootstrap sample. As a result, they are not easily scalable to large-scale healthcare data. In addition, their performance depends on the value of the selection threshold π .

In this paper, we propose a new and simple strategy that leverages the likelihood-based Bayesian machine learning technique, BART, for variable selection with MAR data. By properly using the posterior distributions of the variable importance measures from the BART models on multiple imputed datasets, our proposed approach offers significant computational savings while still delivering excellent performance that is on par with the best performance BI-BART and BI-XGB can achieve with the optimal π . Furthermore, in addition to the commonly used performance metrics for variable selection [9, 11, 19], we propose a way in which each method's ability to select important variables with incomplete data can be judged on the basis of out-of-sample predictive performance via the cross-validated area under the curve (AUC). Finally, we apply our proposed methods to re-analyze the Study of Women's Health Across the Nation (SWAN) data [19] and to identify predictors of 3-year incidence of metabolic syndrome among middle-aged women.

Methods

Proposed methods

Our proposed method uses BART. BART is a tree-based Bayesian probability model [21], offering an additional

advantage over the algorithm-based machine learning methods of providing proper representations of uncertainty intervals via the posterior [22–25]. A regression or classification tree approximates the covariate-outcome relationship by recursive binary partitioning of the predictor space. The tree consists of the tree structure and all the decision rules sending a variable either left or right and leading down to a bottom node. Each of the bottom nodes represents the mean response of the observations falling in that node [14]. For a binary outcome, BART uses probit regression and a sum of trees models,

$$\Pr(Y = 1 | X = \mathbf{x}) = \Phi(f(\mathbf{x})) = \Phi\left(\sum_{j=1}^m g(\mathbf{x}; \mathcal{T}_j, \mathcal{M}_j)\right), \quad (1)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function, $\sum_{j=1}^m g(\mathbf{x}; \mathcal{T}_j, \mathcal{M}_j)$ is a sum of trees model, and $(\mathcal{T}_j, \mathcal{M}_j)$ are tree parameters. To avoid overfitting and limit the contribution of each $(\mathcal{T}_j, \mathcal{M}_j)$, a regularizing prior is put on the parameters, and the posterior is computed using Markov chain Monte Carlo (MCMC). Details of the BART model can be found in Chipman et al. [21].

Here we leverage the underlying Bayesian probability model of the BART technique. We impute the missing data with random forest and use Rubin’s rule [26, 27] to combine the estimates and variances of the variable inclusion proportion (VIP) – the proportion of times each predictor is chosen as a splitting rule divided by the total number of splitting rules appearing in the model – of each predictor provided by the BART model across multiply imputed datasets for variable selection. We refer to our proposed method as RR-BART, where “RR” stands for Rubin’s rule. The key steps of RR-BART are:

- 1 Impute the data M times using the `missForest` technique [28]; fit a BART model to each of M imputed datasets and draw P Markov chain Monte Carlo (MCMC) posterior samples of the VIP for a predictor X_k , $k = 1, \dots, K$. We then have a distribution of VIP_{kmp} , $m = 1, \dots, M$, $p = 1, \dots, P$.
- 2 Calculate the average VIP across P posteriors and M imputed datasets for each predictor $\overline{VIP}_{k..}$, and identify the variable $X_{k'}$ with the minimum value, $\overline{VIP}_{k'..} = \min_{k=1, \dots, K} (\overline{VIP}_{k..})$. If $\overline{VIP}_{k'..} > \frac{1}{2K}$, then stop the algorithm and no variable selection will be performed. Otherwise, calculate the distance between each VIP score and the minimum of average posterior mean VIPs, $\Delta_{kmp} = VIP_{kmp} - \overline{VIP}_{k'..}$ for $k = 1, \dots, K$, $m = 1, \dots, M$, $p = 1, \dots, P$.
- 3 Apply the Rubin’s rule [26, 27] to the distribution of Δ_{kmp} . The overall mean and total variance of the distance score for predictor X_k are calculated as follows:

$$\begin{aligned} \bar{Q}_k &= \sum_{m,p} \Delta_{kmp} / MP, \\ \text{Within-imputation variance } W_k &= \frac{1}{M} \sum_m \text{Var}(\bar{\Delta}_{km.}), \\ \text{Between-imputation variance } B_k &= \frac{1}{M-1} \sum_m (\bar{\Delta}_{km.} - \bar{\Delta}_{k..})^2, \\ \text{Total variance } T_k &= W_k + \left(1 + \frac{1}{M}\right) B_k, \end{aligned}$$

where $\text{Var}(\bar{\Delta}_{km.})$ is the variance among $\{\Delta_{kmp}, p = 1, \dots, P\}$ divided by the sample size n , $\bar{\Delta}_{km.} = \frac{1}{P} \sum_p \Delta_{kmp}$, and $\bar{\Delta}_{k..} = \frac{1}{MP} \sum_m \sum_p \Delta_{kmp}$. The $1 - \alpha$ confidence interval for variable k ’s average distance score is $\bar{Q}_k \pm t_{df, 1-\alpha} \sqrt{T_k}$, where the degrees of freedom for the t -distribution is $df = (M - 1) / ((B_k + B_k/M) / T_k)^2$ [26].

- 4 Select variables with the $1 - \alpha$ confidence intervals that do not contain zero.

The key idea of this approach is to properly characterize the probability distribution of the VIP measure for each predictor from multiple imputed datasets, and then identify the “important” predictors that have significantly larger VIPs. As pointed out by anonymous reviewers, this approach implies that there exists at least one irrelevant predictor so that variable selection is needed. In the extreme scenario where all candidate predictors are useful, the minimum of average posterior mean VIPs, $\overline{VIP}_{k'..}$, will be considerably far away from zero, then variable selection will be deemed as unnecessary and thus all predictors will be selected (step 2). In step 4, bigger values of α will lead to more selected variables and smaller α will result in less variables selected. We used $\alpha = 0.05$ as recommended in previous work [29]. In step 1, we use the flexible imputation method `missForest`. The `missForest` proceeds as follows. Initial guesses are made for the missing values (e.g., mean or mode), and variables are sorted in the ascending order of missingness proportions. The variables with missing data are in turn, in the descending order of missingness proportions, regressed via Random Forest on other variables. The missing values are imputed/updated using the fitted Random Forest model as the prediction model. This process is repeated until a stopping criterion is met or the maximum number of iterations is reached. Detailed description of `missForest` is provided in [Supplementary Section 1](#) and in Stekhoven and Bühlmann [28].

Comparison methods

BI-BART Hu et al. [19] proposed BI-BART, which combines bootstrap imputation with a permutation based variable selection method using BART, developed by Bleich et al. [11] The term “BI” refers to bootstrap imputation. The permutation based BART variable selection method uses the VIP. The response vector is permuted a number

of times (e.g, 100) and the BART model is fitted to each of the permuted response vectors and the original predictors; the VIPs computed from each of BART model fits constitute the null distributions. A variable is selected if its VIP from the BART model fitted to the unpermuted response is larger than the $1 - \alpha$ quantile of its own null distribution. The BI-BART method proceeds with the following three steps: (1) generate B bootstrap data sets and impute missing data using `missForest`, (2) perform variable selection using BART on each bootstrap data set, (3) select the final set of variables if they are selected in at least πB datasets, where π is a fraction threshold between 0 and 1 for selecting a predictor.

BI-XGB BI-XGB combines bootstrap imputation with XGBoost for variable selection among incomplete data. At the core of the XGBoost method is gradient boosting [20]. Boosting is a process in which a weak learner is boosted into a strong learner. The idea of gradient boosting is to build a model via additive functions that minimizes the loss function (exponential loss for classification), which measures the difference between the predicted and observed outcomes [30]. XGBoost uses a gradient tree boosting model with each additive function being a decision tree mapping an observation to a terminal node. In addition, XGBoost uses shrinkage and column subsampling to further prevent overfitting [20]. Shrinkage scales newly added weights by a factor after each step of tree boosting and the column subsampling selects a random subset of predictors for split in each step of tree boosting. Variable selection via XGBoost on fully observed data is carried out in a recursive feature elimination procedure [19], using the variable importance score provided by the XGBoost model. Among incomplete data, BI-XGB carries out variable selection in three steps as described above for BI-BART, with permutation based BART method in step (2) replaced with XGBoost based variable selection method.

MIA-BART and MIA-XGB As tree-based methods, both BART and XGBoost offer a technique, missingness incorporated in attributes (MIA) [20, 29], to handle missing covariate data by treating the missingness as a value in its own right in the splitting rule. The MIA algorithm chooses one of the following three rules for all variables for splitting X and all splitting values x_c : (1) if X is observed and $X \leq x_c$, send this observation left; otherwise, send this observation right. If X is missing, send this observation left, (2) If X is observed and $X \leq x_c$, send this observation left; otherwise, send this observation right. If X is missing, send this observation right, (3) If X is missing, send this observation left; if it is observed, regardless of its value, send this observation right. We refer to work by Kapelner and Bleich [29] for a

detailed description of the MIA procedure. As comparison methods, we implement MIA within the BART model and within the XGBoost model while performing variable selection: MIA-BART and MIA-XGB. Because MIA cannot handle missing outcome data, we implement this technique on two versions of data: (i) only cases with complete outcome data; (ii) data with imputed outcomes.

Performance metrics

To stay consistent with the literature [9, 11, 19] so that our methods and results can be compared with previous work in similar contexts, we assess the performance of variable selection using four metrics: (i) precision, the proportion of truly useful predictors among all selected predictors, (ii) recall, the proportion of truly useful variables selected among all useful variables, (iii) $F_1 = \frac{2 \text{ precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$, the harmonic mean of precision and recall. The F_1 score balances between avoiding selecting irrelevant predictors (precision) and identifying the full set of useful predictors (recall), and (iv) Type I error, the mean of the probabilities that a method will incorrectly select each of the noise predictors. Note that “precision” and “positive predictive value (PPV)” are sometimes used interchangeably; and the same goes for “recall” and “sensitivity”. Type I error is sometimes referred to as “false positive” in the machine learning literature. These performance metrics will be calculated among 250 replications for larger sample sizes $n = 1000, 5000$ and among 1000 replications for smaller sample sizes $n = 300, 650$.

In addition to the four metrics, we propose to use the cross-validated AUC to evaluate the prediction performance of each model with selected variables for incomplete data. This additional metric will be useful to distinguish between the methods in the situation where it remains unclear which method is able to select the most relevant predictors based on the four metrics described above. The AUC evaluation is constituted of the following three steps:

- (1) randomly split the data into two halves, and then implement each of the methods using half of the data to select variables,
- (2) impute the other half of the data (single imputation) and record the AUC of each model with the selected variables. For BART based methods, calculate the AUC using BART, and for XGB based methods, calculate the AUC using XGBoost, and
- (3) repeat steps (1) and (2) 100 times to get the distribution of the AUCs.

Simulation of incomplete data

We adopt the simulation settings used in Hu et al. [19] to accurately mimic realistic missingness problems and to impartially compare methods in similar contexts. We

use the multivariate amputation approach [31] to generate missing data scenarios with desired missingness percentages and the missing data mechanisms. The multivariate amputation procedure first randomly divides the complete data into a certain number (can be one) of subsets, each allowing the specification of any missing data pattern. Then the weighted sum scores are calculated for individuals in each subset to amputate the data. For example, the weighted sum score for X_3 and individual i can take the form, $wss_{x_3,i} = x_{1i}w_1 + x_{2i}w_2$, which attaches a non-zero weight w_1 to x_{1i} and w_2 to x_{2i} to suggest that the missingness in X_3 depends on X_1 and X_2 . A logistic distribution function [32] is then applied on the weighted sum scores to compute the missingness probability, which is used to determine whether the data point becomes missing or not.

We consider simulation scenarios that represent the data structures commonly observed in health studies: (i) sample sizes ranging from small to large: $n = 300, 650, 1000, 5000$, (ii) 10 useful predictors that are truly related to the responses, X_1, \dots, X_{10} , and 10, 20, 40 noise predictors, where X_1 and $X_2 \stackrel{i.i.d.}{\sim} \text{Bern}(0.5)$, X_3, X_4 and $X_5 \stackrel{i.i.d.}{\sim} N(0, 1)$, $X_6 \stackrel{i.i.d.}{\sim} \text{Gamma}(4, 6)$, and X_7, X_8, X_9, X_{10} were designed to have missing values under the MAR mechanism; a half of noise predictors are simulated from $N(0, 1)$, and the other half from $\text{Bern}(0.5)$, and (iii) 20% missingness in Y with 40% overall missingness, and 40% missingness in Y with 60% overall missingness. There are a total of 24 scenarios considered, including 4 sample sizes \times 3 ratios of useful versus noise predictors \times 2 missingness proportions. We additionally investigate how each method performs when there are no noise predictors for $n = 1000$ and two missingness proportions. The outcomes are generated from a model with arbitrary data complexity: (i) discrete predictors with strong (x_1) and moderate (x_2) associations; (ii) both linear and nonlinear forms of continuous predictors; (iii) nonadditive effects (x_4x_9). The outcome model is specified as:

$$\begin{aligned} \Pr(y = 1 | x_1, \dots, x_{10}) = & \text{logit}^{-1}(-2.7 + 1.8x_1 + 0.5x_2 \\ & + 1.1x_3 - 0.4e^{x_5} - 0.4(x_6 - 3.5)^2 + 0.3(x_7 - 1)^3 \\ & + 1.1x_8 - 1.1x_{10} + 5 \sin(0.1\pi x_4x_9) - 0.4x_5x_{10}^2 \\ & + 0.4x_3^2x_8). \end{aligned}$$

Following generating the full data, the multivariate amputation approach was used to amputate X_7, X_8, X_9, X_{10} and Y under the MAR mechanism to create the desired missingness characteristics. Detailed simulation set-up appears in [Supplementary Section 2](#).

Because we generated four predictors, X_7, X_8, X_9 and X_{10} from the normal distribution with the mean depending on other predictors, correlations among covariates

X_3, \dots, X_{10} were established. The correlations range from -0.034 to 0.413 . [Supplementary Table 1](#) shows the Pearson correlations among the predictors X_3, \dots, X_{10} . An anonymous reviewer pointed out that we can measure the strength of MAR by first taking a missing indicator variable for each variable with missing values and regressing it on the other variables related to it being missing, and then estimating the AUC of the regression model. By this approach, the strength of missingness in X_7, X_8, X_9, X_{10} and Y are strong MAR with the AUC ranging from .72 to .92; see [Supplementary Table 2](#).

Case study

We demonstrate and compare the methods by addressing the emerging variable selection problem studied in Hu et al. [19] using the SWAN data. The SWAN was a multicenter, longitudinal study in the U.S. with the goal of understanding women's health across the menopause transition. The SWAN study enrolled 3305 women aged between 42 and 52 in 1996-1997 and followed them up to 2018 annually. The emerging research question is the identification of important risk factors for 3-year incidence of metabolic syndrome. Metabolic syndrome is a cluster of conditions that occur together and has been shown to increase the risk of heart disease, stroke and type 2 diabetes [33]. Identifying important predictors of metabolic syndrome among middle-aged women would be valuable to forming preventive interventions and reducing risks or threats to women's health. This important research question has been less studied in the literature for women during their middle years.

We use the analysis dataset described in Hu et al. [19], which included 2313 women who did not have metabolic syndrome at enrollment. Among the 2313 women, 251 (10.9%) developed metabolic syndrome within three years of enrollment, 1240 (53.6%) did not, and the remaining 822 (35.5%) had missing outcomes. Sixty candidate predictors (29 continuous variables and 31 discrete variables) were selected based on the previous literature [33–36] on risk factors for metabolic syndrome, including the demographics, daily life behaviour, dietary habits, sleep habits, medications, menopausal status and related factors and mental status among others. A detailed list of names and definitions of the 60 candidate variables is provided in [Supplementary Table 12](#). Among the 60 candidate predictors 49 had missing values, and the amount of missing data in these variables ranges from 0.1% to 27.1%. Only 763 (33.0%) participants had fully observed data.

Results

Throughout, we used `missForest` [28] to impute missing data for all methods considered. `missForest` is a flexible imputation algorithm, which recasts the missing data problem as a prediction problem and uses the ran-

dom forest [37] for prediction. The missing values are imputed in turn for each variable from a fitted forest regressing that variable on all other variables [38]. As suggested by Tang et al. [38] and Hu et al. [19], `missForest` performs better than MICE [39] across various missing data settings. All variables (predictor and response variables) available to the analyst were included in the imputation model without variable selection or specification of functional forms. Details of the imputation appear in [Supplementary Section 1](#).

Simulation results

Table 1 summarizes the variable selection results of each method when performed on the fully observed data, incomplete data and complete cases only, for $n = 1000$, 40 noise predictors, and 30% and 60% overall missingness. It is apparent that the performance of bootstrap based methods on incomplete data depends on the threshold π , the optimal value of which varies by methods. By comparison, the performance of RR-BART does not depend on any additional parameters. Even with a large proportion of missingness, RR-BART achieves good variable selection performances (on the bases of AUC, precision, recall, F_1 and Type I error) that are similar to the best performances BI-BART or BI-XGB could achieve with the optimal threshold values of π . Both BART and XGBoost had substantially deteriorated performance when only complete cases were used. Both MIA based methods had subpar performances (only slightly better than those of the complete-case analyses), with the missing outcomes either imputed or excluded.

[Supplementary Table 3–6](#) summarize simulation results for different sample sizes $n = 300, 650, 1000, 5000$ and for $\alpha = 0.01, 0.05, 0.1$. When the sample size is small ($n = 300, 650$), no methods provided satisfactory performance; in the presence of missing data, the proposed method could still recover the performance achievable on fully observed data. The findings are in agreement with the previous literature [19]. When the sample size is large ($n = 1000, 5000$), all three methods RR-BART, BI-BART and BI-XGB, delivered good performance, and comparatively better performance with a smaller missingness proportion and more noise predictors. When the sample size increased from $n = 1000$ to $n = 5000$, the performance gain was only moderate, suggesting that $n = 1000$ is sufficiently large of a sample size for implementing the proposed method. Bigger values of α tend to lead more selected variables, thus increased Type I error and recall. As suggested by an anonymous reviewer, we also included a “baseline” method, referred to as RR-BART (median), by which predictors with the posterior mean VIP exceeding the median value of the VIPs are selected. The performance from the baseline method is comparatively lower than the proposed RR-BART method.

Figure 1 further highlights that the AUCs of BI-BART and BI-XGB depend on the selection threshold π , with the highest AUC achieved at $\pi = 0.1$ for BI-BART and $\pi = 0.3$ for BI-XGB. The proposed RR-BART boasts the same highest AUC without the need to specify a value for π . The performance curves of the other four metrics are shown in [Supplementary Fig. 1](#), which convey the same message.

A perusal of Fig. 2 suggests that RR-BART has a substantially higher power for selecting a discrete variable of even moderate effect size (X_2) than BI-BART and BI-XGB; meanwhile, maintains comparable power for identifying the continuous variables, even in complex conditions of nonadditivity and nonlinearity.

In the extreme case where all of 10 predictors were useful, the average minimum VIP score was around 0.08, which is closer to 1/10 than is to 0. Note that the VIPs of all predictors sum to one. [Supplementary Table 7](#) presents the distribution of the minimums of posterior mean VIPs across 250 data replications. The simulation results of this extreme scenario are provided in [Table 2](#) and [Supplementary Table 8](#). By definition, for *all* methods, the precision is one and the Type I error becomes not applicable. There was a substantial drop in the F_1 score for both BART and XGBoost on fully observed data. This is because the comparatively less important variables were not selected, leading to a lower recall and in turn a lower F_1 score. Turning to the situations in which missing covariate and outcome data are present. When the missingness proportion is smaller (30% overall missingness), the proposed RR-BART delivered similar performance, judged by F_1 and recall, as BI-BART and BI-XGB. When the missingness proportion is higher (60% overall missingness), RR-BART had a lower recall and F_1 score than the two bootstrap imputation based methods. Because the minimum VIP was considerably far away from 0, RR-BART step 2 deemed variable selection as unnecessary and thus selected all predictors. This modified algorithm, RR-BART (all selected), produced near-perfect performance.

[Supplementary Table 9](#) summarizes the numbers of times RR-BART performed better than, worse than or equal to BI-BART, across all simulation configurations. Overall, the two methods produced very similar performance. The small Monte Carlo errors of the variable VIP scores (see [Supplementary Table 10](#)) indicate that the numbers of simulation replications are sufficiently large.

To assess whether the normality assumption underlying Rubin’s rule is satisfied, [Supplementary Fig. 2](#) plots the posterior distributions of the VIPs for the 10 useful predictors, for the simulation scenario with $n = 1000$, 40 noise predictors, and 60% overall missingness proportion. It is reasonable to assume that the VIP scores are normally distributed. In addition, we also explored the

Table 1 Simulation results for each variable selection approach performed on the fully observed data and among incomplete data. For bootstrap imputation based methods on incomplete data, we show results corresponding to different threshold values of π . The optimal value of π leading to the highest F_1 score is $\pi = 0.1$ for BI-BART and $\pi = 0.3$ for BI-XGB. The sample size is $n = 1000$. The number of useful predictors is 10 and the number of noise predictors is 40. Two missingness proportions were considered: 40% missingness in Y and 60% overall missingness; 20% missingness in Y and 40% overall missingness. The performance measures were computed across 250 data replications

	AUC	Precision	Recall	F_1	Type I error
Fully observed data					
BART	0.92 (0.88, 0.96)	1.00	0.87	0.93	0.00
XGBoost	0.88 (0.84, 0.92)	0.93	0.81	0.86	0.02
Incomplete data: 40% missingness in Y and 60% overall missingness					
RR-BART	0.82 (0.78, 0.86)	0.87	0.80	0.83	0.01
BI-BART $\pi = 0.1$	0.83 (0.78, 0.88)	0.87	0.82	0.85	0.03
BI-BART $\pi = 0.2$	0.81 (0.76, 0.86)	0.97	0.71	0.82	0.01
BI-BART $\pi = 0.3$	0.77 (0.72, 0.82)	0.99	0.63	0.77	0.00
BI-BART $\pi = 0.4$	0.73 (0.68, 0.78)	1.00	0.55	0.71	0.00
BI-BART $\pi = 0.5$	0.67 (0.62, 0.72)	1.00	0.48	0.65	0.00
BI-BART $\pi = 0.6$	0.59 (0.54, 0.64)	1.00	0.40	0.57	0.00
BI-BART $\pi = 0.7$	0.49 (0.44, 0.54)	1.00	0.31	0.47	0.00
BI-BART $\pi = 0.8$	0.41 (0.36, 0.46)	1.00	0.23	0.38	0.00
BI-BART $\pi = 0.9$	0.33 (0.28, 0.38)	1.00	0.14	0.25	0.00
BI-BART $\pi = 1.0$	0.25 (0.20, 0.30)	1.00	0.07	0.13	0.00
BI-XGB $\pi = 0.1$	0.58 (0.53, 0.63)	0.40	0.88	0.55	0.36
BI-XGB $\pi = 0.2$	0.74 (0.69, 0.79)	0.66	0.85	0.75	0.15
BI-XGB $\pi = 0.3$	0.82 (0.77, 0.87)	0.83	0.83	0.83	0.03
BI-XGB $\pi = 0.4$	0.76 (0.71, 0.81)	0.96	0.72	0.82	0.01
BI-XGB $\pi = 0.5$	0.70 (0.65, 0.75)	0.99	0.63	0.77	0.00
BI-XGB $\pi = 0.6$	0.64 (0.59, 0.69)	1.00	0.54	0.70	0.00
BI-XGB $\pi = 0.7$	0.59 (0.54, 0.64)	1.00	0.41	0.58	0.00
BI-XGB $\pi = 0.8$	0.47 (0.42, 0.52)	1.00	0.29	0.44	0.00
BI-XGB $\pi = 0.9$	0.35 (0.30, 0.40)	1.00	0.17	0.29	0.00
BI-XGB $\pi = 1.0$	0.28 (0.23, 0.33)	1.00	0.04	0.09	0.00
MIA-BART (Impute missing Y)	0.75 (0.71, 0.79)	0.80	0.75	0.77	0.04
MIA-BART (Exclude missing Y)	0.72 (0.66, 0.78)	0.78	0.70	0.74	0.05
MIA-XGB (Impute missing Y)	0.74 (0.70, 0.78)	0.81	0.73	0.77	0.04
MIA-XGB (Exclude missing Y)	0.71 (0.65, 0.77)	0.75	0.71	0.73	0.08
BART Complete cases	0.70 (0.63, 0.77)	0.90	0.60	0.72	0.03
XGBoost Complete cases	0.73 (0.66, 0.80)	0.90	0.68	0.77	0.04
Incomplete data: 20% missingness in Y and 30% overall missingness					
RR-BART	0.86 (0.82, 0.90)	0.91	0.84	0.87	0.02
BI-BART $\pi = 0.1$	0.87 (0.83, 0.91)	0.91	0.87	0.89	0.01
BI-BART $\pi = 0.2$	0.85 (0.80, 0.90)	0.99	0.76	0.85	0.02
BI-BART $\pi = 0.3$	0.82 (0.77, 0.87)	1.00	0.68	0.82	0.01
BI-BART $\pi = 0.4$	0.77 (0.72, 0.82)	1.00	0.59	0.76	0.01
BI-BART $\pi = 0.5$	0.72 (0.67, 0.77)	1.00	0.54	0.70	0.00
BI-BART $\pi = 0.6$	0.63 (0.58, 0.68)	1.00	0.46	0.61	0.01
BI-BART $\pi = 0.7$	0.53 (0.48, 0.58)	1.00	0.36	0.51	0.00

Table 1 Simulation results for each variable selection approach performed on the fully observed data and among incomplete data. For bootstrap imputation based methods on incomplete data, we show results corresponding to different threshold values of π . The optimal value of π leading to the highest F_1 score is $\pi = 0.1$ for BI-BART and $\pi = 0.3$ for BI-XGB. The sample size is $n = 1000$. The number of useful predictors is 10 and the number of noise predictors is 40. Two missingness proportions were considered: 40% missingness in Y and 60% overall missingness; 20% missingness in Y and 40% overall missingness. The performance measures were computed across 250 data replications (Continued)

	AUC	Precision	Recall	F_1	Type I error
BI-BART $\pi = 0.8$	0.44 (0.39, 0.49)	1.00	0.28	0.42	0.00
BI-BART $\pi = 0.9$	0.38 (0.33, 0.43)	1.00	0.18	0.30	0.00
BI-BART $\pi = 1.0$	0.29 (0.24, 0.34)	1.00	0.12	0.18	0.00
BI-XGB $\pi = 0.1$	0.60 (0.55, 0.65)	0.44	0.90	0.58	0.30
BI-XGB $\pi = 0.2$	0.76 (0.71, 0.81)	0.69	0.87	0.77	0.11
BI-XGB $\pi = 0.3$	0.84 (0.79, 0.89)	0.86	0.85	0.85	0.02
BI-XGB $\pi = 0.4$	0.78 (0.73, 0.83)	0.99	0.75	0.84	0.01
BI-XGB $\pi = 0.5$	0.73 (0.68, 0.78)	1.00	0.65	0.79	0.00
BI-XGB $\pi = 0.6$	0.67 (0.62, 0.72)	1.00	0.57	0.73	0.00
BI-XGB $\pi = 0.7$	0.62 (0.57, 0.67)	1.00	0.44	0.61	0.00
BI-XGB $\pi = 0.8$	0.49 (0.44, 0.54)	1.00	0.32	0.47	0.00
BI-XGB $\pi = 0.9$	0.38 (0.32, 0.42)	1.00	0.20	0.33	0.00
BI-XGB $\pi = 1.0$	0.31 (0.25, 0.35)	1.00	0.08	0.14	0.00
MIA-BART (Impute missing Y)	0.78 (0.74, 0.82)	0.83	0.77	0.79	0.03
MIA-BART (Exclude missing Y)	0.76 (0.70, 0.82)	0.81	0.74	0.78	0.04
MIA-XGB (Impute missing Y)	0.77 (0.73, 0.82)	0.83	0.76	0.79	0.03
MIA-XGB (Exclude missing Y)	0.73 (0.67, 0.79)	0.78	0.73	0.75	0.06
BART Complete cases	0.73 (0.67, 0.79)	0.92	0.64	0.75	0.03
XGBoost Complete cases	0.76 (0.70, 0.82)	0.93	0.71	0.80	0.03

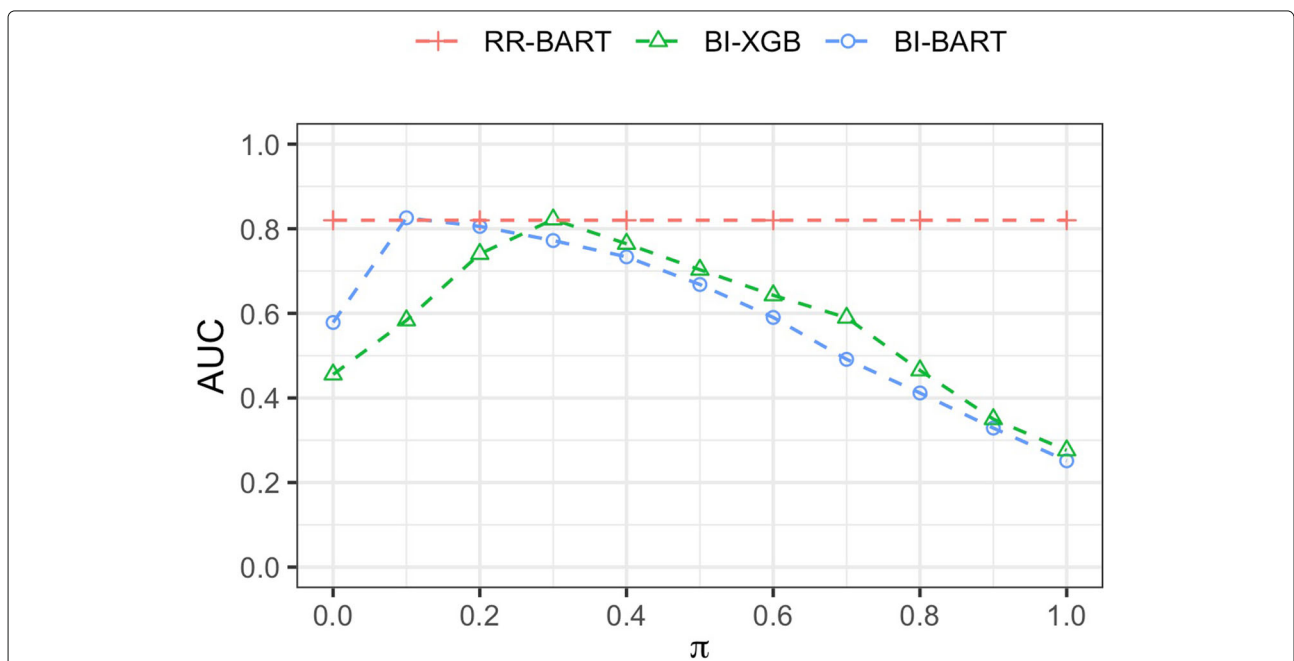


Fig. 1 The mean cross-validated AUC, averaged across 250 data replications, for each of three methods: RR-BART, BI-BART and BI-XGB. The mean AUC for bootstrap imputation based methods BI-BART and BI-XGB varies by the threshold value of π . `missForest` was used for imputation. The sample size $n = 1000$. The proportion of missingness is 40% in the outcome Y and is 60% overall

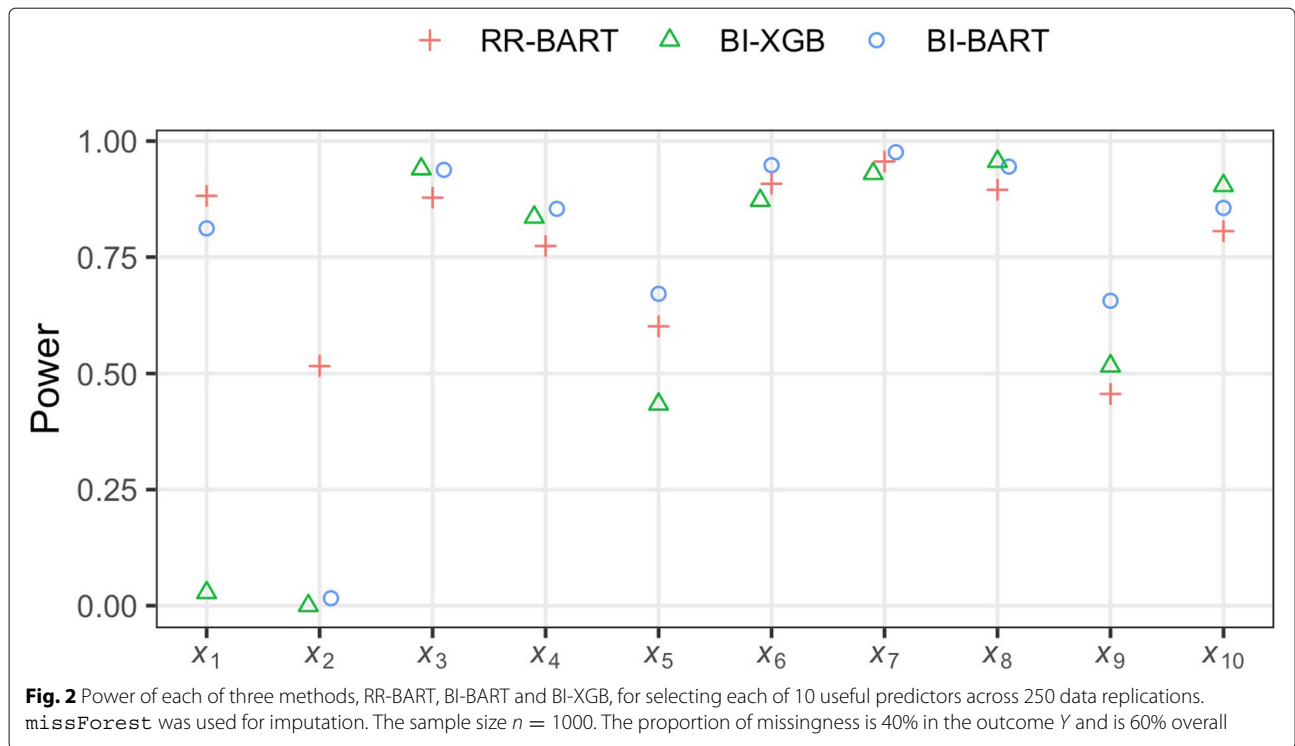


Table 2 Simulation results for the setting in which there are 10 useful predictors and no noise variables. For bootstrap imputation methods on incomplete data, we show results corresponding to the best threshold values of π based on F_1 . The sample size $n = 1000$. Two missingness proportions were considered: 40% missingness in Y and 60% overall missingness; 20% missingness in Y and 40% overall missingness. The performance measures were computed across 250 data replications

	AUC	Precision	Recall	F_1	Type I error
Fully observed data					
BART	0.74 (0.68, 0.80)	1.00	0.62	0.70	NA
XGBoost	0.75 (0.69, 0.81)	1.00	0.61	0.69	NA
Incomplete data: 40% missingness in Y and 60% overall missingness					
RR-BART	0.73 (0.67, 0.79)	1.00	0.36	0.48	NA
RR-BART (all selected)	0.97 (0.95, 0.99)	1.00	1.00	1.00	NA
BI-BART $\pi = 0.1$	0.73 (0.67, 0.79)	1.00	0.38	0.50	NA
BI-XGB $\pi = 0.2$	0.79 (0.73, 0.85)	1.00	0.54	0.64	NA
BART Complete cases	0.50 (0.43, 0.57)	1.00	0.15	0.35	NA
XGBoost Complete cases	0.55 (0.48, 0.62)	1.00	0.18	0.38	NA
MIA-BART (Impute missing Y)	0.66 (0.60, 0.72)	1.00	0.31	0.42	NA
MIA-BART (Exclude missing Y)	0.63 (0.56, 0.70)	1.00	0.25	0.40	NA
MIA-XGB (Impute missing Y)	0.73 (0.67, 0.69)	1.00	0.50	0.59	NA
MIA-XGB (Exclude missing Y)	0.70 (0.62, 0.77)	1.00	0.46	0.55	NA
Incomplete data: 20% missingness in Y and 30% overall missingness					
RR-BART	0.77 (0.72, 0.82)	1.00	0.51	0.67	NA
RR-BART (all selected)	0.98 (0.96, 0.99)	1.00	1.00	1.00	NA
BI-BART $\pi = 0.1$	0.75 (0.70, 0.80)	1.00	0.50	0.69	NA
BI-XGB $\pi = 0.2$	0.80 (0.75, 0.85)	1.00	0.52	0.70	NA
MIA-BART (Impute missing Y)	0.70 (0.65, 0.75)	1.00	0.46	0.60	NA
MIA-BART (Exclude missing Y)	0.67 (0.61, 0.73)	1.00	0.43	0.57	NA
MIA-XGB (Impute missing Y)	0.70 (0.65, 0.75)	1.00	0.46	0.64	NA
MIA-XGB (Exclude missing Y)	0.67 (0.61, 0.73)	1.00	0.42	0.61	NA
BART Complete cases	0.54 (0.49, 0.59)	1.00	0.16	0.39	NA
XGBoost Complete cases	0.56 (0.51, 0.61)	1.00	0.20	0.40	NA

strategy, recommended by Zou and Reiter [40] and Hu et al. [41], which obtains the posterior inferences for the variable VIPs by pooling posterior samples across model fits arising from the multiple datasets. [Supplementary Table 11](#) shows that this strategy produced similar variable selection results as the proposed RR-BART that uses Rubin's rule.

Re-analysis of SWAN data

Table 3 lists the selected important predictors of 3-year incidence of metabolic syndrome by RR-BART, BI-BART $\pi = 0.1$ and BI-XGB $\pi = 0.3$, using the SWAN data. Note that we used the optimal threshold values of π for BI-BART and BI-XGB that were suggested in our simulation study. Both RR-BART and BI-BART identified 17 predictors, among them 15 were the same; 11 predictors were selected by all three methods. In addition to the common risk factors for metabolic syndrome such as the diastolic blood pressure (DIABP), systolic blood pressure (SYSBP), lipoprotein(a) (LPA), and triglycerides (TRI-

GRES), RR-BART was able to identify predictors such as tissue plasminogen activator (TPA) and apolipoprotein A-1 (APOAES) that were less studied in the literature for women's health. The selection of these two predictors has substantial justification, as levels of TPA antigen and low apolipoprotein A-1 (APOAES) were found to be associated with insulin resistance, which was involved in the pathogenesis of impaired fasting glucose [36, 42].

We further evaluated how well the selected variables predict the 3-year incidence of metabolic syndrome. A useful measure to assess the performance of risk prediction models is the validation plot [43]. A perfect risk prediction model would produce a 45° line, with the predicted risk perfectly aligns with the observed event probability. Figure 3 compares the prediction performance of the BART model and the XGBoost model with the predictors selected by RR-BART, BI-BART and BI-XGB. When implemented on the SWAN data, all three models performed well and produced similar AUCs. The validation lines show that RR-BART and BI-BART had better prediction performance for higher risk individuals.

Table 3 Variable selection results by each of 3 methods, with the best imputation method and threshold value of π for BI-BART and BI-XGB suggested in simulations. BART was used with `missForest` and $\pi = 0.1$, XGB with `missForest` and $\pi = 0.3$. Definitions of the variable names appear in Web Table 1. RR-BART and BI-BART both selected 17 variables, and BI-XGB selected 16 variables

Variables	RR-BART	BI-BART	BI-XGB
TRIGRES	Yes	Yes	Yes
SYSBP	Yes	Yes	Yes
LPA	Yes	Yes	Yes
DIABP	Yes	Yes	Yes
WAIST	Yes	Yes	Yes
INSULIN	Yes	No	Yes
LUCRES	Yes	No	No
APOAES	Yes	Yes	Yes
EDUCATION	Yes	No	Yes
BP	Yes	Yes	No
WHRATIO	Yes	Yes	Yes
BMI	Yes	Yes	Yes
RACE	Yes	Yes	No
TPA	Yes	Yes	Yes
DTTLIN	Yes	Yes	No
SHBG	Yes	Yes	Yes
RESTLES	Yes	Yes	No
GLUCOSE	No	Yes	Yes
T	No	Yes	No
E2AVE	No	No	Yes
PAI1	No	No	Yes

Discussion

We investigate strategies for variable selection with missing covariate and outcome data in large-scale health studies. Prior work has shown that combining bootstrap imputation with tree-based variable selection methods has good operating characteristics with respect to selecting the most relevant predictors [19]. The computational cost of this method, however, can be high in large datasets, given that both bootstrap and recursive nonparametric modeling are needed. We propose an inference-based method that leverages the Bayesian machine learning model, BART, and uses Rubin's rule to combine the estimates and variances of the VIPs of each predictor provided by the BART models across multiply imputed datasets for variable selection. In addition, we implement MIA within BART and XGB, which accommodate missing covariates by modifying the binary splitting rules.

Our simulation study suggests that the proposed method RR-BART can achieve the best performance, with respect to AUC, F_1 score, precision, recall and type I error, delivered by the BI-BART or BI-XGB method with the optimal selection threshold value of π . RR-BART can reasonably well recover the performance a BART or XGBoost model can achieve on the fully observed data, in situations where the proportion of missingness is large and in complex conditions of nonlinearity and nonadditivity. Furthermore, it has been shown that tree-based methods tend to have a lower power in detecting a discrete predictor [19]. RR-BART has demonstrated a strong ability of identifying discrete variables, even those which have only moderate effect sizes. In general, using MIA with BART or XGB does not provide satisfactory performance of vari-

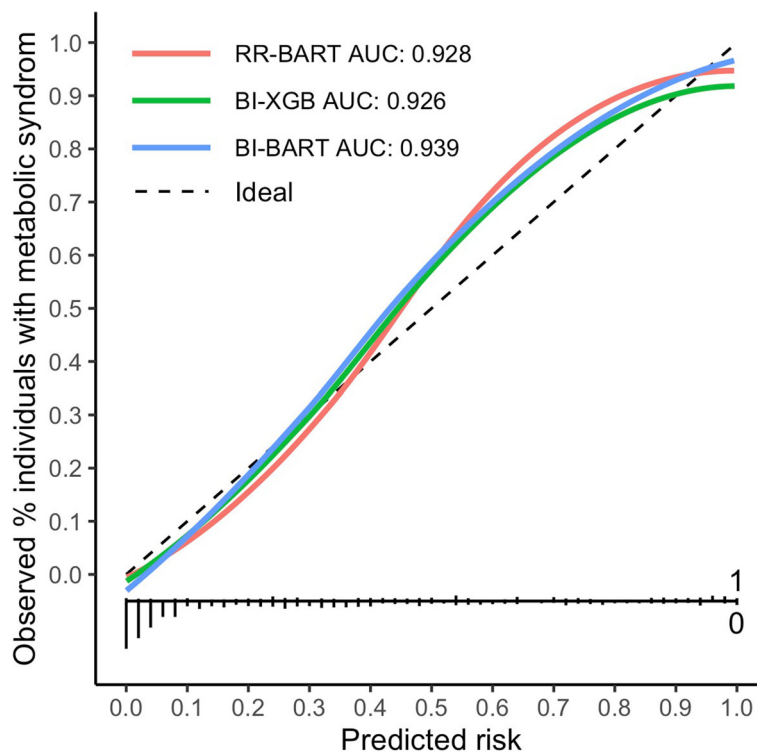


Fig. 3 Validation plot of predicted probabilities of 3-year incidence of metabolic syndrome among middle-aged women using the SWAN data. The risk prediction models were the BART and XGBoost models with predictors selected via RR-BART, BI-BART and BI-XGB. `missForest` was used for imputation

able selection on incomplete data. Moreover, RR-BART recognizes the situation that variable selection may not be needed if the minimum of average posterior mean VIPs is considerably far away from zero.

The computational savings offered by RR-BART are substantial. All simulations were run in R on an iMAC with a 4 GHz Intel Core i7 processor. On a dataset of size $n = 1000$ with 50 predictors, each RR-BART took 4 minutes to run, while each BI-XGB implementation took 7 minutes and each BI-BART took about 55 minutes to run. More importantly, unlike the bootstrap imputation based methods, RR-BART does not require a selection threshold π , whose optimal value varies by methods and simulation settings, and is not known a priori. The performance of RR-BART does depend on the parameter α , for which a default value of 0.05 is recommended.

Although our proposed method RR-BART provides promising performance for variable selection in the presence of MAR covariate and outcome data, it is possible to further improve the method for the big- n -small- p situation, which is frequently encountered in health registry data. Another possible research avenue is to derive an alternative nonparametric measure of variable importance rather than the VIP that reflects the impact of inclusion or deletion of predictors on the model prediction accuracy in the presence of missing data [44].

Finally, extending the methods to accommodate MNAR data could be a worthwhile contribution.

Conclusion

The proposed variable selection method, RR-BART, is both computationally efficient and practically useful in identifying important predictors using large-scale health-care databases with missing data.

Abbreviations

MAR: Missing at random; MCAR: Missing completely at random; MNAR: Missing not at random; BART: Bayesian additive regression trees; XGBoost: Extreme gradient boosting; BI-BART: Bootstrap imputation with Bayesian additive regression trees; BI-XGB: Bootstrap imputation with extreme gradient boosting; RR-BART: Rubin's rule implemented with Bayesian additive regression trees; AUC: Area under the curve; MCMC: Markov Chain Monte Carlo; VIP: Variable inclusion proportion; MIA: Missingness incorporated in attributes; SWAN: Study of Women's Health Across the Nation; DIABP: Diastolic blood pressure; SYSBP: Systolic blood pressure; LPA: Lipoprotein(a); TRIGRES: Triglycerides; TPA: Plasminogen activator; APOAES: Apolipoprotein A-1

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12874-022-01608-7>.

Additional file 1: Web-based supplementary materials. **Supplementary Section 1.** Random Forest based Imputation Algorithm - `missForest`. **Supplementary Section 2.** Simulation Setup. **Supplementary Section 3.** Additional tables (Supplementary Table 1–12) and figures (Supplementary Figure 1–2) for simulation study and case study.

Acknowledgements

Not available.

Authors' contributions

All authors have read and approved the manuscript. JL wrote the final version of the statistical codes, drafted the first version of the manuscript, and reviewed the final version of the paper. LH developed the methods and conceived the simulation study, assembled the team, supervised coding, and finalized the manuscript. CH worked on extracting and cleaning the SWAN data, wrote the first version of the statistical codes and reviewed the final version of the paper. JJ validated the final version of the codes for simulations and data analyses, and reviewed the final version of the paper. SL worked on interpreting the results and on finalizing the manuscript. UG worked on interpreting the results, deciding on its presentation and reviewed the final version of the manuscript.

Funding

This work was supported in part by award ME_2017C3_9041 from the Patient-Centered Outcomes Research Institute, and by grants R21CA245855 and P30CA196521-01 from the National Cancer Institute.

Availability of data and materials

Codes for implementing all methods and replicating simulations are available in the GitHub page of the corresponding author <https://github.com/liangyuanhu/RR-BART>. Deidentified SWAN data are publicly available at <https://www.swanstudy.org/swan-research/data-access>.

Declarations

Ethics approval and consent to participate

All methods were performed in accordance with the relevant guidelines and regulations. All methods implementation, simulations and statistical analyses were performed using R Statistical Software (version 4.0.4; R Foundation for Statistical Computing, Vienna, Austria). Our study is not human subjects research as it consists of only simulated datasets and de-identified data. Because the SWAN data used in our case study are completely de-identified, our institution's institutional review board does not require an approval for research involving SWAN data.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Population Health Science and Policy, Icahn School of Medicine at Mount Sinai, 1425 Madison Ave, 10029 New York, USA.

²Department of Biostatistics and Epidemiology, Rutgers University, 683 Hoes Lane West, 08854 Piscataway, USA. ³Primary Research Solution LLC., 115 W 18th St, 10011 New York, USA.

Received: 4 November 2021 Accepted: 19 April 2022

Published online: 04 May 2022

References

- George EI. The variable selection problem. *J Am Stat Assoc.* 2000;95(452):1304–08.
- Little RJ, D'Agostino R, Cohen ML, Dickersin K, Emerson SS, Farrar JT, Frangakis C, Hogan JW, Molenberghs G, Murphy SA, et al. The prevention and treatment of missing data in clinical trials. *N Engl J Med.* 2012;367(14):1355–60.
- Sterne JA, White IR, Carlin JB, Spratt M, Royston P, Kenward MG, Wood AM, Carpenter JR. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. *BMJ.* 2009;338.
- Little RJ, Rubin DB. *Statistical Analysis with Missing Data*, 3rd edn. New York: Wiley; 2019.
- National Research Council. *The Prevention and Treatment of Missing Data in Clinical Trials*. Washington: The National Academies Press; 2010.
- Hogan JW, Daniels MJ, Hu L. A bayesian perspective on assessing sensitivity to assumptions about unobserved data. In: Molenberghs G, Fitzmaurice G, Kenward MG, Tsiatis A, Verbeke G, editors. *Handbook of Missing Data Methodology*. Boca Raton: CRC Press; 2014. p. 405–34. Chap. 18.
- Hu L, Hogan JW, Mwangi AW, Siika A. Modeling the causal effect of treatment initiation time on survival: Application to HIV/TB co-infection. *Biometrics.* 2018;74(2):703–13.
- Long Q, Johnson BA. Variable selection in the presence of missing data: resampling and imputation. *Biostatistics.* 2015;16(3):596–610.
- Wood AM, White IR, Royston P. How should variable selection be performed with multiply imputed data?. *Stat Med.* 2008;27(17):3227–46.
- Meinshausen N, Bühlmann P. Stability selection. *J R Stat Soc Ser B Stat Methodol.* 2010;72(4):417–73.
- Bleich J, Kapelner A, George EI, Jensen ST. Variable selection for BART: an application to gene regulation. *Ann Appl Stat.* 2014;8(3):1750–81.
- Mazumdar M, Lin J-YJ, Zhang W, Li L, Liu M, Dharmarajan K, Sanderson M, Isola L, Hu L. Comparison of statistical and machine learning models for healthcare cost data: a simulation study motivated by oncology care model (OCM) data. *BMC Health Serv Res.* 2020;20:350.
- Ungaro RC, Hu L, Ji J, Nayar S, Kugathasan S, Denson LA, Hyams J, Dubinsky MC, Sands BE, Cho JH. Machine learning identifies novel blood protein predictors of penetrating and stricturing complications in newly diagnosed paediatric crohn's disease. *Aliment Pharmacol Ther.* 2020;53(2):281–90.
- Hu L, Liu B, Ji J, Li Y. Tree-based machine learning to identify and understand major determinants for stroke at the neighborhood level. *J Am Heart Assoc.* 2020;9(22):016745.
- Hu L, Ji J, Li Y, Liu B, Zhang Y. Quantile regression forests to identify determinants of neighborhood stroke prevalence in 500 cities in the USA: implications for neighborhoods with high prevalence. *J Urban Health.* 2021;98(2):259–70.
- Hu L, Li L, Ji J. Machine learning to identify and understand key factors for provider-patient discussions about smoking. *Prev Med Rep.* 2020;20:101238.
- Hu L, Liu B, Li Y. Ranking sociodemographic, health behavior, prevention, and environmental factors in predicting neighborhood cardiovascular health: a bayesian machine learning approach. *Prev Med.* 2020;141:106240.
- Ji J, Hu L, Liu B, Li Y. Identifying and assessing the impact of key neighborhood-level determinants on geographic variation in stroke: a machine learning and multilevel modeling approach. *BMC Public Health.* 2020;20(1):1–12.
- Hu L, Lin J-YJ, Ji J. Variable selection with missing data in both covariates and outcomes: Imputation and machine learning. *Stat Methods Med Res.* 2021;30(12):2651–71.
- Chen T, Guestrin C. XGBoost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; 2016. p. 785–94.
- Chipman HA, George EI, McCulloch RE. BART: Bayesian additive regression trees. *Ann Appl Stat.* 2010;4(1):266–98.
- Hu L, Lin J, Sigel K, Kale M. Estimating heterogeneous survival treatment effects of lung cancer screening approaches: A causal machine learning analysis. *Ann Epidemiol.* 2021;62:36–42.
- Hu L, Gu C, Lopez M, Ji J, Wisnivesky J. Estimation of causal effects of multiple treatments in observational studies with a binary outcome. *Stat Methods Med Res.* 2020;29(11):287–308.
- Hu L, Gu C. Estimation of causal effects of multiple treatments in healthcare database studies with rare outcomes. *Health Serv Outcome Res Methodol.* 2021;21(3):287–308.
- Hu L, Ji J, Li F. Estimating heterogeneous survival treatment effect in observational data using machine learning. *Stat Med.* 2021;40(21):4691–713.
- Rubin DB. *Multiple Imputation for Nonresponse in Surveys*. New York: Wiley; 2004.
- Hu L, Hogan JW. Causal comparative effectiveness analysis of dynamic continuous-time treatment initiation rules with sparsely measured outcomes and death. *Biometrics.* 2019;75(2):695–707.
- Stekhoven DJ, Bühlmann P. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics.* 2012;28(1):112–18.
- Kapelner A, Bleich J. Prediction with missing data via bayesian additive regression trees. *Can J Stat.* 2015;43(2):224–39.
- Friedman J, Hastie T, Tibshirani R, et al. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *Ann Stat.* 2000;28(2):337–407.

31. Schouten RM, Lugtig P, Vink G. Generating missing values for simulation purposes: a multivariate amputation procedure. *J Stat Comput Simul.* 2018;88(15):2909–30.
32. Van Buuren S. *Flexible Imputation of Missing Data*, 2nd edn. Boca Raton: Chapman & Hall/CRC; 2018.
33. Kazlauskaitė R, Janssen I, Wilson RS, Appelhans BM, Evans DA, Arvanitakis Z, El Khoudary SR, Kravitz HM. Is midlife metabolic syndrome associated with cognitive function change? The Study of Women's Health Across the Nation. *J Clin Endocrinol Metab.* 2020;105(4):1093–105.
34. Han D, Fang X, Su D, Huang L, He M, Zhao D, Zou Y, Zhang R. Dietary Calcium Intake and the Risk of Metabolic Syndrome: A Systematic Review and Meta-Analysis. *Sci Rep.* 2019;9(1):1–7.
35. Janssen I, Powell LH, Crawford S, Lasley B, Sutton-Tyrrell K. Menopause and the metabolic syndrome: the Study of Women's Health Across the Nation. *Arch Intern Med.* 2008;168(14):1568–75.
36. Feng X, Gao X, Yao Z, Xu Y. Low apoA-I is associated with insulin resistance in patients with impaired glucose tolerance: a cross-sectional study. *Lipids Health Dis.* 2017;16(1):1–7.
37. Breiman L. Random forests. *Mach Learn.* 2001;45(1):5–32.
38. Tang F, Ishwaran H. Random forest missing data algorithms. *Stat Anal Data Min: ASA Data Sci J.* 2017;10(6):363–77.
39. Van Buuren S, Boshuizen HC, Knook DL. Multiple imputation of missing blood pressure covariates in survival analysis. *Stat Med.* 1999;18(6):681–94.
40. Zhou X, Reiter JP. A note on Bayesian inference after multiple imputation. *Am Stat.* 2010;64(2):159–63.
41. Hu L, Zou J, Gu C, Ji J, Lopez M, Kale M. A flexible sensitivity analysis approach for unmeasured confounding with multiple treatments and a binary outcome with application to SEER-Medicare lung cancer data. *Ann Appl Stat.* 2022. In press.
42. Rao SS, Disraeli P, McGregor T. Impaired glucose tolerance and impaired fasting glucose. *Am Fam Physician.* 2004;69(8):1961–68.
43. Steyerberg EW, Vickers AJ, Cook NR, Gerds T, Gonen M, Obuchowski N, Pencina MJ, Kattan MW. Assessing the performance of prediction models: a framework for some traditional and novel measures. *Epidemiology.* 2010;21(1):128.
44. Williamson BD, Gilbert PB, Carone M, Simon N. Nonparametric variable importance assessment using machine learning techniques. *Biometrics.* 2021;77(1):9–22.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

