International Journal of Applied Data Science
Volume1, Issue 2, 2025

Understanding Variable Effects from Black Box Prediction: Quantifying Effects in Tree Ensembles Using Partial Dependence

Zaini Sakawi[1] and Mohd Yusof Hussain[2]
1Fakulti Sains Sosial dan Kemanusiaan Universiti Kebangsaan Malaysia.

Article Info

Received: 12-01-2025 **Revised:03-02-2025** **Accepted:17-03-2025** **Published:10-04-2025**

Abstract: Characterizing the impact of one or more factors on a result is often of scientific interest. Although data mining techniques like random forests are adaptable substitutes for traditional parametric models, their lack of interpretability stems from the inability to quantify variable effects in a manner that is substantively relevant. In this study, we offer a partial dependency approach for quantifying variable impacts that averages over the effects of all other variables and yields an estimate that may be understood as the influence on the response for a one unit change in the predictor. Most significantly, the method eliminates issues with model misspecification and implementation difficulties in high-dimensional situations that arise with other methods (e.g., multiple linear regression). We provide a technique for creating a point estimate of this impact magnitude and assess it via simulation. Additionally, we suggest and assess interval estimates using a non-parametric bootstrap. Data used to forecast the age of abalone is used to demonstrate the approach.

Key words*:* Bagging, Bootstrap, Data Mining, Random Forest, Partial Dependence

Introduction

The majority of science focuses on describing how one or more factors affect a result, namely the kind and intensity of the connection. Generally speaking, data mining techniques are characterized by a high degree of flexibility in the number of predictor variables that may be taken into account and how their effects are "modeled," but this flexibility sometimes comes at the price of being able to accurately quantify their effects. When the number of predictors (p) becomes close to the number of observations (n), estimates using traditional parametric models like multiple linear regression become incorrect. On the other hand, data mining techniques like random forests (Breiman, 2001) may be used regardless of the magnitude of p in relation to n and do not need any model definition. However, the incapacity of random forests to provide a substantively relevant characterization of variable effects is a problem. on contrast, a regression coefficient from a multiple linear regression is relevant inasmuch as it indicates the direction and degree of a variable's impact on the response (on the original scale of both variables) while accounting for the influence of other factors. The lack of

One reason they haven't been widely used in scientific applications, where the interpretation of a numerical quantity that characterizes a variable's effect is crucial, may be the lack of techniques to quantify variable effects in random forests in a similar way (Dasgupta et al., 2012; Friedman, 2001). It is often important to be able to evaluate variable effects using effect sizes with a meaning comparable to coefficients from a multiple linear regression, even in research situations where prediction may be the main goal. In this research, we propose an impact size based on the partial dependency technique (Friedman, 2001), which

uses random forests as the foundation for the estimate to overcome issues with model misspecification and implementation in high dimensional situations.

According to Hastie, Tibshirani, and Friedman (2009), tree-based approaches are among the best for data mining applications in a predictive context because they provide a variety of inputs and the capacity to find relevant predictors from a huge pool. In terms of producing precise predictions, random forests (Breiman, 2001) are a significant improvement over earlier tree-based techniques. Combining bagging, also known as bootstrap aggregating (Breiman, 1996), with random predictor selection at each tree split improves prediction accuracy in random forests. By averaging identically distributed trees, bagging lowers variance. This impact is amplified by random feature selection, which further lowers variance by de-correlating the trees (Breiman, 2001; Hastie, Tibshirani, and Friedman, 2009). However, when an ensemble of trees replaces the simplicity of a single tree, this increase in prediction accuracy comes at the expense of interpretability.

To far, a number of methods for interpreting variable effects in random forests have been developed. Using conditional permutation significance (Strobl et al., 2008) or permutation importance (Hastie, Tibshirani, and Friedman, 2009) is one option. Because permutation significance only depicts a variable's predictive capacity in relation to other variables in the forest, it often does not define variable effects in a manner that is substantively significant. Another option is to employ this predictor variable subset in a traditional parametric statistical model that may be used for variable interpretation after identifying the most significant variables (for example, using permutation importance). Given the availability of regression coefficients, this technique does represent variable effects in a substantively relevant manner, but it is not entirely adequate. In particular, unless a very comprehensive variable retention technique is used, it may leave out significant factors and fail to account for the impacts of all variables when estimating the effects of variables in the subset (Strobl, Malley, and Tutz, 2009). Additionally, this method depends on accurate model definition of every variable. Another option is partial dependency (Friedman, 2001), which is defined as the average of the impacts of other factors used to develop the forest and the impact of one or more variables on the response (in their original scale). This approach has only been used for graphical representations so far, despite the fact that it is attractive due to its resemblance to the interpretation of coefficients from a multiple linear regression. In this study, we provide a technique based on the partial dependency approach that may be utilized to produce point estimates and confidence ranges. We suggest and assess via simulation a technique for building a point estimate and non-parametric bootstrap-based interval estimates (Efron, 1987; Efron and Tibshirani, 1993) Data for predicting the age of abalone is used to demonstrate the approach.

1. Methodology

1.1. Random Forests

Random forests can be described in several steps. Adopting the notation of Hastie et al. (2009), the first step is to take B bootstrap samples of size N from the data, where N corresponds to the total number of observations in the data. Grow a tree T_b on each bootstrap sample for $b = 1, \dots, B$. The tree growing process is based on what is typically used for decision trees (Breiman et al., 1984) except that at each split a subset, m, of the total number of variables, p, is selected at random. Furthermore, trees are not pruned, with terminal nodes occurring when the minimum node size is reached. Prediction for an observation x in the context of regression is based on taking the average value on the response for the terminal node in the bth tree where x appears,

In the context of classification, letting $C_b(x)$ be the class prediction of the bth tree, the prediction is conventionally based on:

$$\underset{rf}{C}^{B}(x) = majority\ vote\ \{C_b(x)\}_1^B \quad (2)$$

In the case of two classes, if the relative frequency of the event class is greater than .5 in the terminal node of x for the bth tree, an event would be predicted, whereas if it is less than .5 then no event is predicted. Then, a tree-averaged prediction for an observation is based on whether or not the proportion of trees predicting an event exceeds .5.

1.2. Partial Dependence

Using the notation of Hastie et al. (2009), consider the subvector X_S of $l < p$ for the input predictor variables $X' = (X_1, X_2, \ldots X_p)$ with $S \subset \{1,2, \ldots, p\}$. Further, let C be the complement set with $S \cup C = \{1,2, \ldots, p\}$. X_S is a set of variables whose 'target' effect is of interest, and X_C is the set of all other variables in the data which we seek to average over. Partial dependence is then defined as (Friedman, 2001):

$$f_S(X_S) = \mathrm{E}_{X_C} f(X_S, X_C) = \int f(X_S, X_C) p_C(X_C)\, dX_C \quad (3)$$

with marginal probability density $p_C(X_C) = \int p(X)\ dX_S$, where $p(X)$ is the joint density over all inputs.

The quantity in (3) for a single tree in a forest is estimated by (Friedman, 2001):
$, x_{2C}, \ldots, x_{NC}$ are the values of X_C with N observations. To obtain an estimate of partial dependence for a forest, we would average over the b trees in the forest:$Y^{*} = {}^{1}\sum^{B} f^{*}(X)$
(5)
where $d = 1, \ldots, D$, corresponding to the distinct combination of values taken on by X_S. Generally, X_S can be comprised of multiple variables, each variable having two or more unique values. In this general case D would correspond to the total number of distinct combinations of values across those variables. The utility of considering multiple variables in a partial dependence approach, which might be referred to as 'multivariable partial dependence', is that it allows the possibility of examining interactions among two or more variables while averaging over the
effects of all other inputs. More common however is to consider only a single variable partial dependence (hereafter the only type of partial dependence that is considered). In this case D would correspond simply to the number of distinct values a single variable takes on in a dataset. To obtain a prediction for a particular value, $X_S = x$, from an individual tree that has been grown,

all observations in the dataset are assigned the value of x while keeping the values of all other variables as they are. This synthetic data is then passed through the tree to construct the prediction. To obtain the forest averaged prediction for x, do the same for the remaining B trees and take an average of the predictions. Finally, repeat this process for all other values of X_S occurring in the data. Typically, partial dependence plots have distinct values of X_S plotted on the X-axis and their forest-averaged prediction on the Y-axis.

1.3. An Effect Size Based on Partial Dependence

The general idea of the proposed approach is to use the output produced as part of partial dependence to fit a parametric model, which in turn is used to obtain a point estimate. It is important to note that partial dependence for a single variable will only have the functional form for that variable specified (if that variable is nominal and treated as a categorical input then no restriction on functional form is placed), all the remaining variable inputs which are being averaged over are not subject to parametric restrictions, one of the distinct advantages of using random forests. Of course any misspecification of functional form can be mitigated by fitting a model that corresponds to the form observed in a partial dependence plot.

For continuous outcome data a point estimate is based on the regression coefficient from a weighted least squares regression, with weights corresponding to the frequency with which each distinct value of X_d, $d = 1 \ldots D$, appears in the original dataset. In this case the model is:

$$\hat{Y}_d = X_d\beta + \varepsilon_d \tag{6}$$

with $\hat{\beta} = (X'\Omega^{-1}X)^{-1}X'\Omega^{-1}\hat{Y}$

and $\Omega = diag\ (n_1, \ldots, n_D)$ where n is a frequency

β from above can be interpreted as an average treatment effect (ATE) (Imbens, 2004). This is the effect, in the population, of moving all subjects from being untreated to treated (i.e., in the case of a binary explanatory variable characterized by the presence or absence of treatment). The ATE interpretation follows from the manner in which the outcome, $\hat{Y}_d$, is calculated. That is, predictions are averaged over all observations for each level of an explanatory variable.

Confidence intervals for $\hat{\beta}$ can be based on a nonparametric bootstrap. One option is the percentile method (Efron, and Tibshirani, 1993). A two-sided confidence interval is calculated by using the sorted bootstrap distribution of the estimated effect, $\hat{\beta}^*$, to identify the values of the lower and upper bound given B bootstrap replicates:

A better interval is the bias-corrected and accelerated confidence interval (BC_a), which corrects for bias and skewness in the shape of the bootstrap distribution (Efron, 1987).

To summarize, the proposed algorithm consists of the following:

1) Grow a forest.
2) Estimate partial dependence (for a single variable).
    a. Create D datasets ($d = 1, \ldots, D$). For all observation in the d^{th} dataset only let them take on one value for the variable of interest while keeping values of all other variables unchanged.
    b. Pass the d^{th} dataset through each tree and average the predictions over the trees in the forest.
    c. Repeat Part b for each of the D datasets

3) Construct a point estimate of the proposed effect size by fitting a weighted least squares model with response based on the tree-averaged predicted values obtained in Step 2, the explanatory variable corresponding to the value used to generate each tree-averaged prediction, and weight based on the frequency each value the explanatory variable takes on in the original data.
4) For confidence intervals, repeat Steps 1-3 for as many bootstrap samples as desired

2. Numerical Examples

2.1. Simulation Design

The design of the simulation is based on manipulating the magnitude of the target variable's effect and sample size. We used 10,000 simulated cases per condition to evaluate bias and overall accuracy (root mean square error (RMSE)) of the point estimate, and 1,000 simulated cases per condition to evaluate interval estimates with 1,000 bootstrap replicates per simulated case. Bias was calculated as the parameter minus the estimate, averaged over the number of simulations. The RMSE was calculated by taking the square root of the average squared distance of the estimate from the parameter. Coverage was calculated as the proportion of times the parameter is captured by the 95% bootstrap confidence interval.

Data Generation Process 1

Data generation process (DGP) 1 is based on the following model:
with, X_{ip}^{*} denoting a standardized value corresponding
$X_{i1}, X_{i3}, X_{i4}, X_{i5}, X_{i6}, X_{i7} \sim Discrete\ Uniform\ (0,30)$

$X_{i2}, X_{i8}, X_{i9}, X_{i10}, X_{i11}, X_{i12} \sim Bernoulli\ (.5)$

$\varepsilon_i \sim N(0,1)$

Here X_{i0} equals 1 for all i, therefore β_0 represents the intercept. The remaining variables in the model were independently drawn from the indicated distributions with effects that were zero for all but β_1 and β_2. We consider the situation where $corr(X_{ij}, X_{ij'}) = 0\ for\ j = 1, \ldots, 12$, which is motivated by a desire to initially evaluate and the proposed method in the simplest of contexts, a relatively small number of inputs with no correlation.

Data Generation Process 2

Additional simulations considered the more realistic situation where there is a correlation among the inputs $corr(X_{ij}, X_{ij'}) \neq 0$ for some $X_{ij}s$ combined with higher dimensional noise (adding 30 binary variables with null effects), using the following DGP:

$$Y_{i1} = \beta_0 X_{i0} + \beta_1 X^{*} + \cdots + \beta_{42} X^{*} + \varepsilon_i$$

with, X_{ip}^{*} denoting a standardized value corresponding to:

$X_{i3}, X_{i4}, X_{i5}, X_{i6}, X_{i7} \sim Discrete\ Uniform\ (0,30)$

$X_{i1}, X_{i2}, X_{i8}, X_{i9}, \ldots, X_{i42} \sim Bernoulli\ (.5)$

$\varepsilon_i \sim N(0,1)$
$i42$

with $corr(X_{i2}, X_{i8}) = corr(X_{i1}, X_{i8}) = corr(X_{i1}, X_{i2}) = \rho$ and zero correlation otherwise Values

of ρ were set to 0, .25, .50, or -.33 (this was the maximum negative correlation that could be simulated in this context without obtaining a non-positive definite covariance matrix). We changed $X_{i1} \sim Bernoulli\ (.5)$ from the previous simulation in order to generate data from a multivariate bernoulli distribution with the desired correlation structure. We also added 30 binary noise variables.

Standardization of variables in both DGP 1 & 2 facilitated assigning values to β_1 and β_2 that were comparable. Parameters β_1 and β_2 were set to be equal and may take on values of .050, .331, 1.134 corresponding to small, medium, or large associations on a correlation scale (r= .05, .30, .60), respectively. The choice of simulating from a discrete uniform distribution is based on a desire to mimic the effect of a continuous random variable. Unfortunately, use of a truly continuous random variable would be computationally too expensive in the context of a simulation study (but not necessarily in an applied context) because it would require generating as many datasets as there are distinct values of the continuous random variable, and then passing each through each tree in the forest. The $X_{i1}, \ldots, X_{i12}$ variables were inputs into the random forest for the DGP-1 and $X_{i1}, \ldots, X_{i42}$ for DGP-2. The effect based on partial dependence was calculated using (6) and confidence intervals were constructed using the nonparametric bootstrap described in (7), as well as a bias-corrected and accelerated version of this interval estimator. We considered 3 effects (.050, .331, 1.134) X 5 sample sizes (40, 100, 250, 500, 1000) x 5 correlation structures/inputs (ρ=0/12inputs, ρ=0/42 inputs, ρ=.25/42 inputs, ρ=.50/42 inputs, ρ=-.33/42 inputs) for point estimation, and for interval estimates we evaluated a more limited number of conditions, in particular excluding conditions with N=1000 given the substantial computational burden.

Computation

The DGPs and fitting of random forests was undertaken in R, the latter based on the randomForest package. The DGP for correlated Xs was based on the MultiOrd package, which implements the method of Emrich and Piedmonte (1991). The default methods in the randomForest package were used for all except one parameter, the number of trees. This was set to 300 because we found little evidence of improvement with increasing number of trees. Nonparametric bootstrapping was implemented using the boot package of R for both the percentile and the bias-corrected and accelerated confidence intervals. The number of bootstrap replicates used for the construction of each interval for each simulated cases was 1000. Given the computationally burdensome design, simulating cases was done in parallel using the snowfall package. Evaluation of point estimation accuracy was executed in parallel on a desktop PC with multi-core processor, whereas larger tasks involving evaluation of interval estimates for any one condition was more problematic (1000 simulated cases x 1000 bootstrap replicates= 1,000,000 random forests) and therefore required use of a computing cluster, the Triton Shared Computing Cluster at the University of California-San Diego.

2.2. Results

2.2.1. DGP 1 Results

Our initial simulations manipulated the number of variables considered at each split. Table 1 displays the results of selecting the number of variables considered at each split (mtry) to be chosen based on a search function (tuneRF) that minimizes the out of bag error for each simulated case. Selection of the number of predictors begins with 4 predictors at each split, and increases or decreases the number of predictors by a factor of 2 (or truncated to the maximum number of predictors-12), stopping when the level of improvement in the out of bag error is less than .01. The results in Table 1 are less than ideal. In Figure 1 we illustrate the effect of manual selection of different numbers of predictors to be used at each split on parameter estimation. The figure conveys a clear improvement with increasing mtry, with the best performance occurring when mtry is set at the maximum. For this reason simulations reported hereafter only considered mtry set at its maximum.

Table 2 displays bias and RMSE as a function of the parameter value and sample size for non-null effects. One result that is immediately clear is that any underestimation bias and overall inaccuracy (as indexed by RMSE) is attenuated by increasing the magnitude of the effect (this decreases the relative bias) and increasing sample size. We have found that the pattern of reduced bias and improved accuracy with increasing effect and sample size corresponds to a greater frequency of these non-null variables appearing in the first two splits of a tree (see Appendix A). Therefore, appearance in early splits appears to be an important determinant of obtaining accurate point estimates. In contrast, for parameters with null effects in the model (e.g., $\beta_3 = \beta_8 = 0$), the bias was calculated to be less than .001 irrespective of the sample size and the magnitude of parameters with non-null effects in the model. As expected, for those variables with null effects the RMSE does decrease with increasing sample size (Table 2).

Based on the coverage reported in Table 3 for non-null effects, we observed improved coverage with increasing effect size and sample size, which is not surprising given the aforementioned improvement in parameter estimation with increasing sample size and effect size. Interestingly, the BCa intervals generally provide superior coverage to the percentile method with medium and large effects while the percentile method appears superior with small and null effects. The difference between the two methods generally decreases with increasing sample size, therefore for larger datasets the decision about which method to use may be less important. In practice, we might prefer BCa intervals when, for instance, there is a notable difference between the mean of the bootstrap estimates and the estimate from the full sample. It is worth noting that we did not find evidence of improvement with an increase in the number of bootstrap replicates. For instance, we used 5,000 bootstrap replicates for the N=40, $\beta = .050$ condition, the coverage with BCa intervals for the non-null effects was .758/.590, which is not substantially different from the Table 3 values of .769/.612.

2.2.2. DGP 2 Results

It does not appear that increasing the number of noise variables alone alters the degree of bias in the estimates (cf. Table 4 for ρ=0 to Table 2 entries). A positive correlation among the inputs leads to less bias in non-null effects than when there is no correlation among the inputs, and this effect increases with increasing positive correlation. In a few limited conditions there is a

tendency for positive correlation to lead to an overestimation bias, but the magnitude of the bias is relatively small. In contrast, when there is a negative correlation among the inputs, this leads

to increased underestimation bias relative to when there is no correlation for both non-null effects and null effects. We should note for conditions with a negative correlation, the bias decreases with increasing sample size. Collectively, this suggests that higher dimensional noise has little effect on estimation, but the sign of the correlation among inputs can either improve or worsen estimation of this quantity, with improvements generally resulting with increasing sample size.

2.3. Application to Prediction of Abalone Age

The data for this example originates from the Tasmanian Aquaculture and Fisheries Institute, obtainable from the UCI machine learning repository. One way of regulating the harvesting of abalone is through limiting which abalone can be harvested based on their age. Age of abalone can be determined by counting the rings on its shell. Unfortunately, this requires cutting the shell, which would not be an effective way of regulating its harvesting. Therefore, if the age of the abalone can be determined from other physical measurements that would be preferable. The data contained in the dataset aim to achieve this goal. The input variable used in the prediction of the number of rings (+1.5 is a proxy for Abalone age) included: Sex (Infant, Male, Female), Length, Diameter, Height, Whole Weight, Shucked Weight, Viscera Weight, and Shell Weight. We fit a random forest with rings as the response and the variables described above as inputs. We used 300 trees for the random forest and the maximum mtry (8). We used 5000 bootstraps replicates to construct 95% Percentile and BCa confidence intervals. We report on the effects of two variables, sex and length using the proposed method. R programming syntax for this example is provided as supplementary material.

The partial dependence plot for length is provided in Figure 2. The estimate for length was -.562, comparison of infants to females was -.242, and comparison of males to females was .004. There was notable bias in the estimate of length (i.e., estimate of mean from full sample-mean of bootstrap estimates= -.151), suggesting that BCa would be more appropriate. Similar bias was not present for either of the estimates involving comparison of the sex variable (-0.009 (I) & -0.007 (M)). For length the BCa method gives 95% confidence intervals of (-1.325, 0.264) and percentile (-1.671, 0.067). For the infant effect the BCa method gives 95% confidence intervals of (-0.368, -0.035) and percentile (-0.394, -0.077). For the male effect the BCa method gives 95% confidence intervals of (-0.075, 0.094) and percentile (-0.093, 0.079).

3. Discussion

Quantifying a variable's effect is important for the adoption of random forests in future scientific applications. Our simulations suggest estimates are less biased and more accurate with increasing effect and sample size. An important determinant for obtaining unbiased and accurate results appears to be whether variables with true effects appear in early splits. Confidence intervals constructed using a non-parametric bootstrap appears to be an effective method, one that improves with increasing effect and sample size as well.

There are several limitations associated with the simulation study that could lead to fruitful directions for future research. First, we did not consider dichotomous outcome data. One

implementation of an approach involving such outcome data could involve calculating partial dependence predictions based on the proportion of events in the terminal nodes of each tree averaged over the trees in the forest, as opposed to predictions based on majority vote, which have been shown in some simulations to be less accurate (Malley et al., 2012). If interest only centers on interpreting binary explanatory variables, it would be straightforward to calculate the risk difference, risk ratio, or odds ratio based on probabilities obtained via partial dependence. For explanatory variables with many response options a feasible alternative would be a weighted beta regression model (Ferrari and Cribari-Neto, 2004) with use of a logit link function leading to an odds ratio interpretation of the exponentiated regression coefficients. Yet another outcome type that is often of interest is time to event, and quantification of a variable's effect could be based on a survival probability calculated using partial dependence at a fixed time point (Ishwaran et al., 2008)

There are alternative modeling approaches that could be considered to achieve the same goals as those considered in this paper. One option that has recently been proposed in the data mining literature is counterfactual machines (Dasgupta et al., 2014). In the simple case of a binary explanatory variable, two forests are built, each using only data from observations belonging to each level of the variable. Predictions for each observation then are based on passing each observation through the forest they were not used to construct, thus providing counterfactual inferences regarding what would have happened if the observation belonged to the other level of the explanatory variable. Although this is an appealing approach, it can present challenges in implementation when explanatory variables are not binary. Yet another approach might involve the use of propensity scores in a conventional parametric model (Rosenbaum and Rubin, 1983). It should be noted however that such approaches rely on correct specification of the propensity score model, which may be difficult to achieve. One possibility is use of a tree-based ensemble method, such as random forests, to estimate each observation's probability of receiving "treatment" (Lee, Lessler, and Stuart, 2010). However, there are important practical limitations with the use of propensity scores. One such problem is that going through the process of using propensity scores for each variable that might be of interest (in terms of the effect that variable has on the outcome) in a research study with a large number of explanatory variables can be quite cumbersome. Therefore, the method proposed in this paper might be considered a reasonable alternative.

Several factors may make the implementation of the methods described in this paper challenging. Larger sample sizes, numbers of inputs, and response options for input variables whose effect will be calculated will increase computing time. In particular, bootstrapping confidence intervals in this context can be a computationally expensive proposition. Fortunately, bootstrapping is a task that is easily parallelized and computation can be sped up by parallelizing on a multi-core processor or a computing cluster. We provide R programming syntax for the empirical example that illustrates how this can easily be accomplished with a multi-core processor. Moreover, increasing computing speeds over time will further increase the feasibility of bootstrapping with large datasets. The method proposed in this paper represents a reasonable approach to quantifying the effect in random forests. The method will allow for a wider adoption of random forest methodology, in particular to applications where the nature and strength of a variable's effect on the outcome is of interest.

References

[1]. Breiman, L. (1996). Bagging Predictors. Machine Learning, 24, 123–140.

[2]. Breiman, L. (2001). Random Forests. *Machine Learning,* 45, 5–32.

[3]. Breiman, L., Friedman, J.H., Olshen, R.A., and Stone, C.J. (1984). *Classification and Regression Trees*. New York: Chapman & Hall.

[4]. Dasgupta, A., Szymczak, S., Moore, J.H., Bailey-Wilson, J.E., and Malley, J.D. (2014). Risk estimation using probability machines. *BioData Mining,* 7, 1-17.

[5]. Efron, B. and Tibshirani, R.J. (1993). *An Introduction to the Bootstrap*. London, U.K.: Chapman & Hall.

[6]. Efron, B. (1987). Better Bootstrap Confidence Intervals. *Journal of the American Statistical Association,* 82, 171-200.

[7]. Emrich, L.J., and Piedmonte, M.R. (1991). A method for generating high-dimensional multivariate binary variates. *American Statistician,* 45, 302-304.

[8]. Ferrari, S.L.P. and Cribari-Neto, F. (2004). Beta regression for modeling rates and proportions. *Journal of Applied Statistics,* 31, 799–815.

[9]. Friedman, J. (2001). Greedy Function Approximation: The Gradient Boosting Machine. *Annals of Statistics,* 29, 1189-1232.

[10].Hastie,T., Tibshirani, R., and Friedman, J. (2009).*The Elements of Statistical Learning: Prediction, Inference and Data Mining*. New York: Springer Verlag.

[11].Imbens, G.W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *The Review of Economics and Statistics,* 86, 4–29.

[12].Ishwaran, H., Kogalur, U.B., Blackstone, E.H., Lauer, M.S. (2008). Random survival forests. *Annals of Applied Statistics,* 2, 841-860.

[13].Lee, B.K., Lessler, J., and Stuart, E.A. (2010). Improving propensity score weighting using machine learning. *Statistics in Medicine,* 29, 337-346.

[14].Malley, J.D., Kruppa, J., Dasgupta, A., Malley, K.G., Ziegler, A. **(**2012). Probability machines: consistent probability estimation using nonparametric learning machines. *Methods of Information in Medicine,* 51, 74-81.

[15].Rosenbaum, P.R., and Rubin, D.B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika,* 70, 41–55.

[16].Strobl, C., Boulesteix, A.L., Kneib, T., Augustin, T., Zeileis A. (2008). Conditional