
Prediction Error Reduction Function as a Variable Importance Score

E. Fokoué^{1*}

¹*School of Mathematical Sciences, Rochester Institute of Technology,
98 Lomb Memorial Drive, Rochester, NY 14623, USA.*

Article Type Here

Abstract

This paper introduces and develops a novel variable importance score function in the context of ensemble learning, and demonstrates its appeal both theoretically and empirically. Our proposed score function is simple and more straightforward than its counterpart proposed in the context of random forest, and by avoiding permutations, it is by design computationally more efficient than the random forest variable importance function. Just like the random forest variable importance function, our score handles both regression and classification seamlessly. One of the distinct advantage of our proposed score is the fact that it offers a natural cut off at zero, with all the positive scores indicating importance and significance, while the negative scores are deemed indications of insignificance. An extra advantage of our proposed score lies in the fact it works very well beyond ensemble of trees and can seamlessly be used with any base learners in the random subspace learning context. Our examples, both simulated and real, demonstrate that our proposed score does compete mostly favorably with the random forest score.

Keywords: High-dimensional, Variable Importance, Random Subspace Learning, Out-of-Bag Error, Random Forest, Large p small n , Classification, Regression, Ensemble Learning, Base Learner

2010 Mathematics Subject Classification: 62H30, 62H25.

1 Introduction

Consider a data set $\mathcal{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)\}$ where $\mathbf{x}_i$ is a p -dimensional vector of attributes of potentially different types observable on some input space denoted here by $\mathcal{X}$, and $\mathbf{y}_i$ are the responses taken from $\mathcal{Y}$. We shall consider various scenarios, but mainly the regression scenario with $\mathcal{Y} = \mathbb{R}$ and the classification scenario with $\mathcal{Y} = \{1, 2, \dots, K\}$. We consider the task of building the estimator $\hat{f}(\cdot)$ of the true but unknown underlying f , and seek to build $\hat{f}(\cdot)$ such that the true error (generalization error) is as small as possible. In this context, we shall use the average test error $\text{AVTE}(\cdot)$, as our measure of predictive performance, namely

$$\text{AVTE}(\hat{f}) = \frac{1}{R} \sum_{r=1}^R \left\{ \frac{1}{m} \sum_{j=1}^m \ell(\mathbf{y}_j^{(r)}, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) \right\}, \quad (1.1)$$

where $(\mathbf{x}_j^{(r)}, \mathbf{y}_j^{(r)})$ is the j th observation from the test set at the r th random replication of the split of the data. Throughout this paper, we shall use the zero-one loss (1.2) for all our classification tasks.

$$\ell(\mathbf{y}_j^{(r)}, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) = 1_{\{\mathbf{y}_j^{(r)} \neq \hat{f}^{(r)}(\mathbf{x}_j^{(r)})\}} = \begin{cases} 1 & \text{if } \mathbf{y}_j^{(r)} \neq \hat{f}^{(r)}(\mathbf{x}_j^{(r)}) \\ 0 & \text{otherwise.} \end{cases} \quad (1.2)$$

For regression tasks, we shall use the squared error loss (1.2), namely

$$\ell(\mathbf{y}_j^{(r)}, \hat{f}^{(r)}(\mathbf{x}_j^{(r)})) = (\mathbf{y}_j^{(r)} - \hat{f}^{(r)}(\mathbf{x}_j^{(r)}))^2. \quad (1.3)$$

Besides seeking the optimal predictive estimator of f , we also seek to select the most important (useful) predictor variables as a byproduct of our overall learning scheme. Indeed, while accurate prediction is very important in and of itself, it's often desirable or even crucial in some cases, provide the added description of the importance of the variables involved in the prediction task. The statistical literature is filled with thousands of papers on variable selection and measurement of variable importance. [1] and [2] propose a measure of variable importance for random forest. [8] and [7] have written many interesting contributions to the estimation of variable importance specifically in the context of classification and regression trees and their random forest ensemble learning extension. [3] implements some of those measures of variable importance for tree-based learning in their wonderful R package `caret`. [6] give a nice summary of variable importance and [5] also touches on this subject. [4] proposed a bias-correction improvement to the above measures of variable importance. Unlike all these authors whose work on variable importance measures concentrated solely on tree-based models, we herein propose a measure that goes beyond tree-based methods.

2 Main result

We consider a framework with a p -dimensional input space $\mathcal{X}$ with typical input vector $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_p)^\top$. We also consider building different models with different subsets of the p original variables. Let $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)^\top$ denote the p -dimensional indicator such that

$$\gamma_j = \begin{cases} 1 & \text{if } \mathbf{x}_j \text{ is active in the current model indexed by } \boldsymbol{\gamma} \\ 0 & \text{otherwise.} \end{cases} \quad (2.1)$$

Assume that we are given an ensemble (collection or aggregation) of models, say

$$\mathcal{H} = \{\mathbf{g}(\cdot, \boldsymbol{\gamma}^{(1)}), \mathbf{g}(\cdot, \boldsymbol{\gamma}^{(2)}), \dots, \mathbf{g}(\cdot, \boldsymbol{\gamma}^{(B)})\} \quad (2.2)$$

where $\mathbf{g}(\cdot, \boldsymbol{\gamma}^{(b)})$ denotes the function built with only those variables that are active in the b th model of the ensemble (aggregation), and $\boldsymbol{\gamma}^{(b)} = (\gamma_1^{(b)}, \dots, \gamma_p^{(b)})$ with

$$\gamma_j^{(b)} = \begin{cases} 1 & \text{if } \mathbf{x}_j \text{ is active in the } b\text{-th model of the ensemble} \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

For instance, we may consider a homogeneous ensemble, i.e, an ensemble in which all the functions are of the same family, like the case where all the base learners are multiple linear regression (MLR) models differing by the variables upon which they are built. Consider a score function $\text{score}(\mathbf{g}(\cdot, \boldsymbol{\gamma}^{(b)}))$ used to assess the performance of model indexed by the variables active in $\boldsymbol{\gamma}^{(b)}$. We propose a variable importance score in the form of a function that measures the importance of a variable $\mathbf{x}_j$ in terms of the reduction in average score

$$\text{PERF}(\mathbf{x}_j) = \frac{1}{B} \sum_{b=1}^B \text{score}(\mathbf{g}(\cdot, \boldsymbol{\gamma}^{(b)})) - \frac{1}{B_j} \sum_{b=1}^B \gamma_j^{(b)} \text{score}(\mathbf{g}(\cdot, \boldsymbol{\gamma}^{(b)})) \quad (2.4)$$

where B_j is the number of models containing the variable $\mathbf{x}_j$, specifically $B_j = \sum_{b=1}^B 1_{\{\gamma_j^{(b)}=1\}}$. In words,

$$\text{PERF}(\mathbf{x}_j) = \text{Average score over all models} - \text{Average score over all models with } \mathbf{x}_j$$

Intuitively, $\text{PERF}(\mathbf{x}_j)$ somewhat measures the impact of variable $\mathbf{x}_j$. In the way similar to the approach used by sports writers to decide the MVP on a team or in a league, $\text{PERF}(\mathbf{x}_j)$ looks at the overall performance of the whole ensemble and then for each variable $\mathbf{x}_j$ computes the direction and magnitude of the change to that overall performance of the ensemble brought by its presence in models. *If a variable $\mathbf{x}_j$ is important, then its presence in any model will cause that model to perform better in the sense of having a lower than common average error (score). The average score of all models containing an important variable $\mathbf{x}_j$ should therefore be lower than the overall average score.*

- $|\text{PERF}(\mathbf{x}_j)|$ measures the magnitude of the importance/impact.
- $\text{sign}(\text{PERF}(\mathbf{x}_j))$ measures the direction of the impact.
- If $\text{sign}(\text{PERF}(\mathbf{x}_j)) = +1$ and $|\text{PERF}(\mathbf{x}_j)|$ is relatively large, then $\mathbf{x}_j$ is an important variable.
- Seamlessly applied to large p small n .
- All variables with $\text{PERF}(\mathbf{x}_j) \leq 0$ are unimportant and can be discarded.
- The $\text{PERF}(\cdot)$ score can be used whenever an ensemble $\mathcal{H}$ is available along with a suitable score function for each base learner.
- This works with any base learner and can be adapted to parametric, nonparametric and semi-parametric models and one can imagine ensembles with any base learners as its atoms.
- A great advantage over the traditional variable importance [1], [2] score functions is that the clear cut-off at zero, in the sense that all variables with $\text{PERF}(\mathbf{x}_j) > 0$ are kept and all those variables with $\text{PERF}(\mathbf{x}_j) \leq 0$ are thrown away.

2.1 PERF score via Random Subspace Learning

A natural implementation of $\text{PERF}(\cdot)$ can be done using the ubiquitous bootstrap along with the random subspace learning scheme. The **Out-of-Bag** (oob) error in the bagging or random subspace learning context is a good (in fact excellent) candidate score function, especially when the goal is the selection of variables that lead to the lowest prediction error. The advantage of using oob as the score lies in the fact that the score is obtained as part of building the ensemble in the random subspace learning framework. Consider the training set $\mathcal{D} = \{\mathbf{z}_i = (\mathbf{x}_i^\top, y_i)^\top, i = 1, \dots, n\}$, where $\mathbf{x}_i^\top = (x_{i1}, \dots, x_{ip})$ and $y_i \in \mathcal{Y}$ are realizations of two random variables X and Y respectively. Let $\mathbf{x}_{i,\pi_j} = (x_{i,1}, \dots, x_{i,\pi_j}, \dots, x_{i,d})$. The permutation π_j acts the $|\bar{\mathcal{D}}^{(b)}|$ -dimensional j th column of the out-of-bag data matrix. Essentially, π_j simply permutes the $|\bar{\mathcal{D}}^{(b)}|$ elements of the j th column of the out-of-bag data matrix.

Algorithm 2.1. 1: *procedure* PERF SCORE(B) $\triangleright$ Computing the PERF Score based on B base learners

- 2: Choose a base learner $\hat{g}(\cdot)$ $\triangleright$ e.g.: Trees, MLR
- 3: Choose an estimation method $\triangleright$ e.g.: Recursive Partitioning or OLS
- 4: Initialize all the $\widehat{\text{PERF}}(\mathbf{x}_j)$ and $\widehat{\text{VI}}(\mathbf{x}_j)$ at zero
- 5: **for** $b = 1$ to B **do**
- 6: Draw with replacement from $\mathcal{D}$ a bootstrap sample $\mathcal{D}^{(b)} = \{\mathbf{z}_1^{(b)}, \dots, \mathbf{z}_n^{(b)}\}$
- 7: Draw without replacement from $\{1, \dots, p\}$ a subset $\mathcal{V}^{(b)} = \{j_1^{(b)}, \dots, j_d^{(b)}\}$ of d variables.

8: Form the indicator vector $\gamma^{(b)} = (\gamma_j^{(b)}, \dots, \gamma_p^{(b)})$ with

$$\gamma_j^{(b)} = \begin{cases} 1 & \text{if } j \in \{j_1^{(b)}, \dots, j_d^{(b)}\} \\ 0 & \text{otherwise.} \end{cases}$$

9: Drop unselected variables from $\mathcal{D}^{(b)}$ so that $\mathcal{D}_{\text{sub}}^{(b)}$ is d dimensional

10: Build the b th base learner $\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})$ based on $\mathcal{D}_{\text{sub}}^{(b)}$

11: Compute score of the b th base learner $\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})$ $\triangleright$ e.g. Out-of-bag error

$$\mathbf{s}^{(b)} = \text{score}(\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})) = \frac{1}{|\mathcal{D}^{(b)}|} \sum_{\mathbf{z}_i \notin \mathcal{D}^{(b)}} \ell(y_i, \widehat{\mathbf{g}}(\mathbf{x}_i, \gamma^{(b)}))$$

12: **for** $j \in \mathcal{V}^{(b)}$ **do**

13: Generate the permutation of the j th column of $\bar{\mathcal{D}}^{(b)}$, namely

$$\pi_j$$

14: Compute the permutation impacted score

$$\mathbf{s}_{\pi_j}^{(b)} = \text{score}_{\pi_j}(\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})) = \frac{1}{|\mathcal{D}^{(b)}|} \sum_{\mathbf{z}_i \notin \mathcal{D}^{(b)}} \ell(y_i, \widehat{\mathbf{g}}(\mathbf{x}_i, \pi_j, \gamma^{(b)}))$$

15: Compute the b th instance of the importance of $\mathbf{x}_j$

$$\widehat{\text{VI}}^{(b)}(\mathbf{x}_j) = \mathbf{s}^{(b)} - \mathbf{s}_{\pi_j}^{(b)}$$

16: **end for**

17: **end for**

18: Use the ensemble $\mathcal{H} = \{\widehat{\mathbf{g}}(\cdot, \gamma^{(b)}), b = 1, \dots, B\}$ to form the estimator

$$\widehat{\text{PERF}}(\mathbf{x}_j) = \frac{1}{B} \sum_{b=1}^B \text{score}(\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})) - \frac{1}{B_j} \sum_{b=1}^B \gamma_j^{(b)} \text{score}(\widehat{\mathbf{g}}(\cdot, \gamma^{(b)})) \quad (2.5)$$

$$\widehat{\text{VI}}(\mathbf{x}_j) = \frac{1}{B_j} \sum_{b=1}^B \gamma_j^{(b)} \widehat{\text{VI}}^{(b)}(\mathbf{x}_j) \quad (2.6)$$

19: **end procedure**


3 Computational demonstrations

3.1 Simulated Example

The dataset in this example is simulated data with different scenarios on the level of correlation among the variables, and the ratio n and p . In this particular example, the true function is

$$f(\mathbf{x}) = 1 + 2\mathbf{x}_3 + \mathbf{x}_7 + 3\mathbf{x}_9$$

with $\mathbf{x} \sim \text{MVN}(\mathbf{1}_9, \Sigma_\rho)$ and $\epsilon \sim \text{N}(0, 1)$. The dataset in this example is simulated data with different scenarios on the level of correlation among the variables, and the ratio n and p . Specifically, we simulate data by defining $\rho \in [0, 1]$, then we generate our predictor variables using a multivariate



(a) Permutation-free Variable Importance.

(b) Permutation-based Variable Importance.

Figure 1: Variable score for simulated data with high correlation among the variables in low dimension high sample size setting

normal distribution. Throughout this paper, the multivariate Gaussian density will be denoted by $\phi_p(\mathbf{x}; \boldsymbol{\mu}, \Sigma)$

$$\phi_p(\mathbf{x}; \boldsymbol{\mu}, \Sigma) = \frac{1}{\sqrt{(2\pi)^p |\Sigma|}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\} \quad (3.1)$$


Furthermore, in order to study the effect of the correlation pattern, we simulate the data using a covariance matrix Σ parameterized by τ and ρ and defined by $\tau\Sigma$ where $\Sigma = (\sigma_{ij})$ with $\sigma_{ij} = \rho^{|i-j|}$.

$$\Sigma = \Sigma(\tau, \rho) = \tau \begin{pmatrix} 1 & \rho & \cdots & \rho^{p-2} & \rho^{p-1} \\ \rho & 1 & \rho & \cdots & \rho^{p-2} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \rho^{p-2} & \ddots & \rho & 1 & \rho \\ \rho^{p-1} & \rho^{p-2} & \cdots & \rho & 1 \end{pmatrix}$$

For simplicity however, we use the first Σ with $\tau = 1$ throughout this paper. For the remaining parameters, we use $\rho \in \{0, 0.25, 0.75\}$ and $p \in \{17, 250\}$, with the same $n = 200$.


4 Conclusion and Discussion

We have presented a variable importance score function in the context of ensemble learning. Our proposed score function is simple and more straightforward than its counterpart proposed in the context of random forest, and by avoiding permutations, it is by design computationally more efficient than the random forest variable importance function. Just like the random forest variable importance function, our score handles both regression and classification seamlessly. One of the distinct advantage of our proposed score is the fact that it offers a natural cut off at zero, with




(a) Permutation-free Variable Importance. (b) Permutation-based Variable Importance.

Figure 2: Variable Importance Scores for simulated data with mild correlation among the variables in low dimension high sample size setting




(a) Permutation-free Variable Importance. (b) Permutation-based Variable Importance.

Figure 3: Variable Importance Scores for simulated data with zero correlation among the variables in low dimension high sample size setting




(a) Permutation-free Variable Importance.




(b) Permutation-based Variable Importance.

Figure 4: Variable Importance Scores for the Attitude Data Set, for which $n = 30$ and $p = 6$.



(a) Permutation-free Variable Importance.



(b) Permutation-based Variable Importance.

Figure 5: Variable Importance Scores for the Spam Detection Dataset where $n = 200$ and $p = 7$, and $K = 2$ classes.

-
- [2] L. Breiman. Statistical modeling: The two cultures. *Statistical Science*, 16(3):199–215, August 2001.
 - [3] M. Kuhn. Building predictive models in r using the caret package. *Journal of Statistical Software*, 28(5), 2008.
 - [4] M. Sandri and P. Zuccolotto. A bias correction algorithm for the gini variable importance measure in classification trees. *Journal of Computational and Graphical Statistics*, 17(3):611–628, 2008.
 - [5] Y. S. Shih. Regression trees with unbiased variable selection. *Statistica Sinical*, 12:361–386, 2009.
 - [6] C. Strobl, , T. Hothorn, and A. Zeileis. Party on. *R Journal*, 1(2):14–17, December 2009.
 - [7] C. Strobl, A-L. Boulesteix, T. Kneib, T. Augustin, and A. Zeileis. Conditional variable importance for random forests. *BMC Bioinformatics*, 9(307):1–12, July 2008.
 - [8] C. Strobl, A-L. Boulesteix, A. Zeileis, and T. Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8(25), 2007.
-