THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Bayesian Bootstrap Inference for the ROC Surface

Citation for published version:

Calhau Fernandes Inacio De Carvalho, V, de Carvalho, M & Branscum, A 2018, 'Bayesian Bootstrap Inference for the ROC Surface', *Stat.* <https://doi.org/10.1002/sta4.211>

Digital Object Identifier (DOI):

[10.1002/sta4.211](https://doi.org/10.1002/sta4.211)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Stat

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

ARTICLE TYPE

Bayesian Bootstrap Inference for the ROC Surface

Vanda Inácio de Carvalho^{*1} | Miguel de Carvalho¹ | Adam Branscum²

¹School of Mathematics, University of Edinburgh, Scotland, UK

²College of Public Health and Human Sciences, Oregon State University, Oregon, USA

Correspondence

*Vanda Inácio de Carvalho, JCMB, The King's Buildings, EH9 3FD. Email: Vanda.Inacio@ed.ac.uk

Summary

Accurate diagnosis of disease is of great importance in clinical practice and medical research. The receiver operating characteristic (ROC) surface is a popular tool for evaluating the discriminatory ability of continuous diagnostic test outcomes when there exist three ordered disease classes (e.g., no disease, mild disease, advanced disease). We propose the Bayesian bootstrap, a fully non-parametric method, for conducting inference about the ROC surface and its functionals, such as the volume under the surface (VUS). The proposed method is based on a simple, yet interesting, representation of the ROC surface in terms of placement variables and has the appealing feature of producing point and interval estimates for the ROC surface and its corresponding VUS in a single integrated framework. Results from a simulation study demonstrate the ability of our method to successfully recover the true ROC surface and to produce valid inferences in a variety of complex scenarios. An application to data from the Trail Making Test to assess cognitive impairment in Parkinson's disease patients is provided.

KEYWORDS:

Diagnostic test, Bayesian bootstrap, ROC surface, volume under the surface, trail making test

1 | INTRODUCTION

Evaluating and ranking the performance of medical diagnostic tests is of fundamental importance in health care. Before a test is approved for routine use in practice, its ability to distinguish between different disease stages or conditions must be rigorously evaluated through statistical analysis. The receiver operating characteristic (ROC) curve is a popular tool for evaluating the accuracy of continuous outcome diagnostic tests that classify subjects into two groups: diseased and nondiseased. However, disease progression can be regarded as a dynamic process and, in clinical practice, physicians often face situations that require a decision among three or more diagnostic alternatives. Patients may advance through one or more transitional stages prior to full disease onset, and this is especially true for neurological disorders. For instance, in Section 5 we present an assessment of the discriminatory ability of the Trail Making Test, a widely used test to detect cognitive impairment associated with dementia, to distinguish between Parkinson's disease patients who present normal cognition, mild cognitive impairment, and dementia/severe impairment. As a direct generalisation of ROC curves, the ROC surface has been used to evaluate diagnostic accuracy in ordered three-class classification problems (Nakas & Yiannoutsos 2004; Yang & Carlin 2000); see also Nakas (2014) for an insightful review of three-class ROC surface analysis. The volume under the ROC surface (VUS) has been proposed as a scalar summary measure of diagnostic accuracy, analogous to the area under the ROC curve in the two-class setting.

There is a vast literature on parametric, semiparametric, and nonparametric ROC curve analysis; see Pepe (2003) and Zhou, McClish, and Obuchowski (2011) for an overview. The amount of existing research on ROC surface analysis is, by comparison, limited. Li and Zhou (2009) developed a frequentist nonparametric approach to estimating the ROC surface based on the empirical distribution function and a semiparametric approach that attempts to generalise the parametric (normal) functional form of the ROC surface but that, as pointed out by the authors, relies heavily on the normality assumption. Inácio, Turkman, Nakas, and Alonzo (2011) developed a nonparametric Bayesian method based on finite mixtures of Polya trees to estimate the ROC surface, while Zhang and Li (2011) developed methods for combining multiple biomarkers, Yu (2012) considered rare

diseases and high-throughput data, Li, Zhou, and Fine (2012) considered ROC surface regression analysis, Kang and Tian (2013) developed nonparametric estimators based on kernel methods, and Coolen-Maturi and Coolen (2014) considered a frequentist nonparametric predictive approach. An empirical likelihood approach (Wan 2012) and inverse probability weighting (Zhang & Alonzo 2016) have been used to estimate the volume under the ROC surface, and test statistics were developed by Hong and Cho (2015).

In this paper, we develop a computationally appealing, fully nonparametric estimator of the ROC surface based on the Bayesian bootstrap. The Bayesian bootstrap (Rubin 1981) is the Bayesian analog of the frequentist bootstrap. A key difference is that the Bayesian bootstrap is based on simulation rather than resampling (details about the Bayesian bootstrap are given in Section 2). Our nonparametric estimator of the ROC surface is robust and smooth. Compared to a kernel approach, it does not depend on a smoothing parameter, the choice of which is a nontrivial issue in practice and has a great impact on inference. Our method is hence a widely applicable approach to inference for the ROC surface that can be used for many populations and for a large number of diseases and continuous diagnostic measures. Moreover, point estimates and credible intervals for the ROC surface and its corresponding VUS are obtained in a single integrated framework. Recent developments of flexible Bayesian models that have been successfully applied in medical diagnostic testing research abound (e.g., Branscum, Johnson, and Baron 2013; Branscum, Johnson, Hanson, and Baron 2015; Branscum, Johnson, Hanson, and Gardner 2008; Erkanli, Sung, Jane Costello, and Angold 2006; Hwang and Chen 2015; Inácio de Carvalho and Branscum 2018; Inácio de Carvalho, de Carvalho, and Branscum 2017; Inácio de Carvalho, Jara, Hanson, and de Carvalho 2013; Rodríguez and Martínez 2013; Zhao, Feng, Chen, and Taylor 2016).

The remainder of the paper is organised as follows. In the next section we introduce background material on ROC surfaces and the Bayesian bootstrap. Section 3 presents our novel nonparametric approach to estimating the ROC surface. The performance of our method is assessed in Section 4 using simulated data, Section 5 applies our approach to data from a Trail Making Test to detect cognitive impairment, and Section 6 concludes the paper. The Trail Making Test data and an R function for implementing our method are provided in the Supplementary Materials.

2 | PRELIMINARIES

2.1 | ROC surfaces

We assume that each subject in the population belongs to one of three ordered diagnostic groups (e.g., no disease, mild disease, advanced disease) and that a diagnostic test with continuous scale outcomes is used for classification into one of the three groups. We further assume that the group to which each subject belongs to is known without error due to the existence of a gold standard test. Without loss of generality, we assume that individuals from group 3 tend to have higher test outcomes than individuals in group 2 who tend to have higher test values than group 1 individuals. Let Y_1 , Y_2 , and Y_3 be continuous random variables denoting test outcomes in groups 1, 2, and 3, respectively, with F_1 , F_2 , and F_3 being the corresponding cumulative distribution functions. For any pair of ordered thresholds (c_1, c_2) , with $c_1 < c_2$, the probabilities of correct classification into each group are given by

$$\begin{aligned} p_1(c_1, c_2) &= \Pr(Y_1 \leq c_1) = F_1(c_1), \\ p_2(c_1, c_2) &= \Pr(c_1 < Y_2 \leq c_2) = F_2(c_2) - F_2(c_1), \\ p_3(c_1, c_2) &= \Pr(Y_3 > c_2) = 1 - F_3(c_2). \end{aligned}$$

The ROC surface is then the three-dimensional plot in the unit cube depicting the probabilities of correct classification into each group as the thresholds c_1 and c_2 vary

$$\{(p_1(c_1, c_2), p_2(c_1, c_2), p_3(c_1, c_2)) : c_1 < c_2\} = \{(F_1(c_1), F_2(c_2) - F_2(c_1), 1 - F_3(c_2)) : c_1 < c_2\}.$$

For notational simplicity, hereafter we drop the dependence of p_1 , p_2 , and p_3 on c_1 and c_2 .

By writing $c_1 = F_1^{-1}(p_1)$ and $c_2 = F_3^{-1}(1 - p_3)$, we obtain the functional form of the ROC surface

$$\text{ROCS}(p_1, p_3) = \begin{cases} F_2(F_3^{-1}(1 - p_3)) - F_2(F_1^{-1}(p_1)), & \text{if } F_1^{-1}(p_1) < F_3^{-1}(1 - p_3), \\ 0, & \text{otherwise.} \end{cases}$$

The volume under the ROC surface is a summary measure of the overall diagnostic accuracy and it is defined as

$$\begin{aligned} \text{VUS} &= \int_0^1 \int_0^1 \text{ROCS}(p_1, p_3) dp_3 dp_1 \\ &= \int_0^1 \int_0^{1-F_3(F_1^{-1}(p_1))} \left\{ F_2(F_3^{-1}(1-p_3)) - F_2(F_1^{-1}(p_1)) \right\} dp_3 dp_1 \\ &= \Pr(Y_1 < Y_2 < Y_3). \end{aligned}$$

When the three distributions completely overlap, and thus the test has no discriminatory ability, the VUS takes the value 1/6, whereas a VUS of 1 corresponds to a test that perfectly discriminates between the three groups. Other values of VUS correspond to different degrees of overlap/stochastic ordering between f_1/F_1 , f_2/F_2 , and f_3/F_3 ; the closer the VUS is to 1, the better the classification accuracy (see Figure 1 of the Supplementary Materials).

2.2 | Bayesian bootstrap

The Bayesian bootstrap (BB) was introduced by Rubin (1981) as a Bayesian counterpart of the original bootstrap proposed by Efron (1979) and it is based on simulation rather than resampling. Let $(y_1, \dots, y_n)$ be a random sample from an unknown distribution F and suppose that the parameter of interest is F itself, which is represented as

$$F(z) = \sum_{i=1}^n \omega_i I(y_i \leq z),$$

where $I(\cdot)$ denotes the indicator function, ω_i is the weight associated to observation y_i , with $\omega_i \geq 0$ and $\sum_{i=1}^n \omega_i = 1$. In the classical nonparametric bootstrap, inference about F is obtained by repeatedly generating bootstrap samples, where each bootstrap sample is drawn with replacement from the data. In the b th bootstrap replicate, $F^{(b)}$ is computed as

$$F^{(b)}(z) = \sum_{i=1}^n \omega_i^{(b)} I(y_i \leq z),$$

where $\omega_i^{(b)}$ is the proportion of times y_i appears in the b th bootstrap sample, with $\omega_i^{(b)}$ taking values on the discrete set $\{0, 1/n, \dots, n/n\}$. For the Bayesian bootstrap, the weights are considered unknown and their posterior distribution is derived. Rubin (1981) used a diffuse prior, $\prod_{i=1}^n \omega_i^{-1}$, which when combined with the multinomial likelihood for the data (assuming no ties), results in a Dirichlet($n; 1, \dots, 1$) distribution for the posterior distribution of the weights. Thus, the weights in the BB are smoother than those from the classical bootstrap. Note that in the BB the data are regarded as fixed, so we do not resample from it. The BB has connections to the Dirichlet Process (Ferguson 1973); specifically, it can be regarded as a non-informative version of the Dirichlet Process (Gasparini 1995, Theorem 2). For a further explanation of the different views of the BB we refer the reader to Kim, Lee, and Kim (2005, p. 971).

3 | PROPOSED ESTIMATOR

Our estimator extends to the three class-setting the method proposed by Gu, Ghosal, and Roy (2008) for the ROC curve. It is motivated by a simple, yet interesting and computationally appealing representation of the ROC surface that is based on the notion of a placement variable (Pepe 2003, Chapter 5). A placement variable is simply a standardisation of test outcomes with respect to a reference population. Consider the following two variables

$$U_1 = F_1(Y_2), \quad U_3 = 1 - F_3(Y_2),$$

with U_1 being the proportion of class 1 subjects with test outcomes less than or equal to Y_2 and U_3 being the proportion of class 3 subjects with test outcomes greater than Y_2 . Here, group 2 is the reference group. The variables U_1 and U_3 quantify the degree of separation of the test outcome distributions in the three groups of patients. Specifically, U_1 quantifies the degree of separation between the test outcomes in groups 1 and 2, whereas U_3 quantifies the degree of separation between groups 2 and 3. For instance, if the test outcomes in the three groups are highly separated, the placement of most group 2 subjects is at the upper tail of the group 1 distribution and at the lower tail of the group 3 distribution, so that most group 2 subjects will have large U_1 and U_3 values. On the other hand, when the three distributions of test outcomes completely overlap, both U_1 and U_3 will have a Uniform(0, 1) distribution.

Interestingly, the ROC surface is the difference between the survival distribution of U_3 and the cumulative distribution of U_1 . Specifically, if

$F_1^{-1}(p_1) < F_3^{-1}(1 - p_3)$, we have

$$\begin{aligned} \text{ROCS}(p_1, p_3) &= F_2(F_3^{-1}(1 - p_3)) - F_2(F_1^{-1}(p_1)) \\ &= \Pr(Y_2 \leq F_3^{-1}(1 - p_3)) - \Pr(Y_2 \leq F_1^{-1}(p_1)) \\ &= \Pr(1 - F_3(Y_2) > p_3) - \Pr(F_1(Y_2) \leq p_1) \\ &= \Pr(U_3 > p_3) - \Pr(U_1 \leq p_1). \end{aligned} \quad (1)$$

Let $(y_{11}, \dots, y_{1n_1})$, $(y_{21}, \dots, y_{2n_2})$, and $(y_{31}, \dots, y_{3n_3})$ be independent (within and between groups) samples of size n_1 , n_2 , and n_3 from groups 1, 2, and 3, respectively. The result in (1) provides the rationale for the following computational algorithm. Let B denote the number of iterations.

Step 1: Computation of placement variables based on the BB.

For $b = 1, \dots, B$, let

$$U_{1j}^{(b)} = F_1^{(b)}(y_{2j}) = \sum_{i=1}^{n_1} v_{1i}^{(b)} I(y_{1i} \leq y_{2j}),$$

and

$$U_{3j}^{(b)} = 1 - F_3^{(b)}(y_{2j}) = \sum_{\ell=1}^{n_3} v_{3\ell}^{(b)} I(y_{3\ell} > y_{2j}),$$

where $j = 1, \dots, n_2$, $(v_{11}^{(b)}, \dots, v_{1n_1}^{(b)}) \sim \text{Dirichlet}(n_1; 1, \dots, 1)$, and $(v_{31}^{(b)}, \dots, v_{3n_3}^{(b)}) \sim \text{Dirichlet}(n_3; 1, \dots, 1)$.

Step 2: Generate a random realisation of the ROC surface.

Based on (1), generate a realisation of $\text{ROCS}^{(b)}(p_1, p_3)$, the difference between the survival function of $(U_{31}^{(b)}, \dots, U_{3n_2}^{(b)})$ and the distribution function of $(U_{11}^{(b)}, \dots, U_{1n_2}^{(b)})$, i.e.,

$$\text{ROCS}^{(b)}(p_1, p_3) = \begin{cases} \sum_{j=1}^{n_2} w_{3j}^{(b)} I(U_{3j}^{(b)} > p_3) - \sum_{j=1}^{n_2} w_{1j}^{(b)} I(U_{1j}^{(b)} \leq p_1), & \text{if } \sum_{j=1}^{n_2} w_{3j}^{(b)} I(U_{3j}^{(b)} > p_3) > \sum_{j=1}^{n_2} w_{1j}^{(b)} I(U_{1j}^{(b)} \leq p_1), \\ 0, & \text{otherwise,} \end{cases}$$

where p_1 and p_3 span grids over $[0, 1]$, $(w_{11}^{(b)}, \dots, w_{1n_2}^{(b)}) \sim \text{Dirichlet}(n_2; 1, \dots, 1)$, and $(w_{31}^{(b)}, \dots, w_{3n_2}^{(b)}) \sim \text{Dirichlet}(n_2; 1, \dots, 1)$.

The BB estimate of the ROC surface, denoted as $\widehat{\text{ROCS}}(p_1, p_3)$ is obtained by averaging over the ensemble of ROC surfaces $\{\text{ROCS}^{(1)}(p_1, p_3), \dots, \text{ROCS}^{(B)}(p_1, p_3)\}$, that is,

$$\widehat{\text{ROCS}}(p_1, p_3) = \frac{1}{B} \sum_{b=1}^B \text{ROCS}^{(b)}(p_1, p_3).$$

Similarly, the posterior mean for the VUS can be computed as

$$\widehat{\text{VUS}} = \frac{1}{B} \sum_{b=1}^B \text{VUS}^{(b)}, \quad \text{VUS}^{(b)} = \int_0^1 \int_0^1 \text{ROCS}^{(b)}(p_1, p_3) dp_3 dp_1,$$

where the integral is approximated numerically. A credible interval for the VUS can be obtained from the percentiles of the ensemble $(\text{VUS}^{(1)}, \dots, \text{VUS}^{(B)})$.

4 | SIMULATION STUDY

A simulation study was conducted to evaluate the performance of our approach to conduct inference about the ROC surface and its associated VUS.

4.1 | Simulation scenarios

We considered four scenarios as listed in Table 1. Scenario 1 corresponds to the case where test outcomes from the three groups follow normal distributions. In Scenario 2, data from the three groups follow non-normal distributions from the same family, namely the family of gamma distributions. In Scenario 3, test outcomes arise from different distributional families. Lastly, Scenario 4 considers mixtures of distributions for test outcome data, a setting that is common in practice.

4.2 | Models

For our BB estimator we only need to specify the number of BB iterates B , which we set equal to 2000. For the grid of values for p_1 and p_3 , the probabilities of correct classification in group 1 and 3, respectively, we used 50 equidistant points on $[0, 1]$.

We compared the performance of our nonparametric BB estimator against its main nonparametric competitors, namely, the frequentist kernel estimator and the empirical estimator. The empirical method simply estimates F_d by its empirical distribution function. For the kernel estimator, the cumulative distribution function in each group is estimated as

$$\hat{F}_d(y) = \frac{1}{n_d} \sum_{i=1}^{n_d} \Phi\left(\frac{y - y_{di}}{h_d}\right), \quad d \in \{1, 2, 3\},$$

where $\Phi(\cdot)$ stands for the standard normal distribution function. For the bandwidth h_d , which controls the amount of smoothing, we considered two options. The first option was

$$h_d = 0.9 \min\{\text{SD}(\mathbf{y}_d), \text{IQR}(\mathbf{y}_d)/1.34\} n_d^{-0.2}, \quad (2)$$

where $\text{SD}(\mathbf{y}_d)$ and $\text{IQR}(\mathbf{y}_d)$ are the standard deviation and interquartile range, respectively, of $\mathbf{y}_d = (y_{d1}, \dots, y_{dn_d})$. This is the default choice in the R statistical software (R Core Team 2017) and it is implemented in the function `bw.nrd0`. It is well-known (e.g., Wand & Jones 1994, p. 61) that this ‘rule’, although providing reasonable bandwidths for non-normal data, is ‘optimal’ when the data distribution is normal. For this reason, we have also considered a bandwidth selected by least squares cross-validation (Wand & Jones 1994, Chapter 3), which is implemented in R by the function `bw.ucv`. Estimation of VUS for the empirical and kernel approaches used the following closed form expressions (see e.g., Kang & Tian 2013):

$$\begin{aligned} \widehat{\text{VUS}}_e &= \frac{1}{n_1 n_2 n_3} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{\ell=1}^{n_3} I(y_{1i} < y_{2j} < y_{3\ell}), \\ \widehat{\text{VUS}}_k &= \frac{1}{n_1 n_2 n_3} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{\ell=1}^{n_3} \Phi\left(\frac{y_{2j} - y_{1i}}{\sqrt{h_1^2 + h_2^2}}\right) \Phi\left(\frac{y_{3\ell} - y_{2j}}{\sqrt{h_2^2 + h_3^2}}\right), \end{aligned}$$

where VUS_e and VUS_k stand, respectively, for the empirical and kernel VUS.

For Scenario 1 (where test outcomes in each group follow a normal distribution), we also included a comparison to a model involving independent parametric normal distributions, in order to assess the efficiency of our nonparametric estimator in this context.

4.3 | Results

For each of the four scenarios described in Section 4.1, 300 datasets were generated using sample sizes of $(n_1, n_2, n_3) = (50, 50, 50)$, $(n_1, n_2, n_3) = (100, 100, 100)$, and $(n_1, n_2, n_3) = (200, 200, 200)$. The discrepancy between the estimated and true ROC surface was measured by the empirical mean squared error

$$\begin{aligned} \text{EMSE} &= \frac{1}{n_{p1}} \frac{1}{n_{p3}} \sum_{u_1=1}^{n_{p1}} \sum_{u_3=1}^{n_{p3}} \left\{ \widehat{\text{ROCS}}(p_{1u_1}, p_{3u_3}) - \text{ROCS}(p_{1u_1}, p_{3u_3}) \right\}^2 \\ &\approx \int_0^1 \int_0^1 \left\{ \widehat{\text{ROCS}}(p_1, p_3) - \text{ROCS}(p_1, p_3) \right\}^2, \end{aligned}$$

where $n_{p1} = n_{p3} = 50$. The estimated VUS and the EMSEs for Scenarios 1–4 are presented in Figures 1 to 4. Specifically, for each scenario and sample size considered, we present a boxplot of the VUS estimates (along with the true value) and EMSEs produced by each method. In addition, the estimated ROC surfaces, along with the true surfaces, are shown in Figures 2–5 of the Supplementary Materials. In Scenario 1, we can appreciate a minor loss in efficiency of our BB estimator, which is a small price to pay for the benefit of the robustness that leads to accurate data driven estimates under increasingly complex scenarios (Figures 2 to 4). The BB estimator outperformed, in terms of the empirical mean squared error, the empirical estimator for most of the scenarios, especially for the sample size $(n_1, n_2, n_3) = (50, 50, 50)$. The BB estimator was on par with the kernel estimator, as measured by the EMSE (although the kernel method produces slightly more biased estimates), with the additional advantage of not needing to select a smoothing parameter and of providing simultaneously both point and interval estimates. As expected, uncertainty associated with our BB estimator decreased as the sample size increased. Lastly, the frequentist coverage of the 95% credible intervals for the VUS are presented in Table 2. We found the coverages to be between 0.95 and 1, which shows the validity of our inferences.

5 | APPLICATION

The Trail Making Test (TMT) is a neuropsychological test that provides information about visual search speed, scanning, speed of processing, as well as, executive functioning. The TMT test is commonly used as a diagnostic test of cognitive impairment associated with dementia. The TMT comprises two parts, both consisting of 25 circles distributed over a sheet of paper or on a computer screen. In Part A, the circles are numbered from 1 to 25 and patients are tasked with connecting them in a sequential order (1–2–3, etc). In part B, the patient alternates between numbers and letters

(1-A-2-B, etc). The goal is to finish both parts of the test as quickly as possible, and completion times are used as the primary performance metrics. While Part A is primarily used to assess cognitive processing speed, Part B is used to examine executive functioning.

We analysed TMT Part A completion times for 245 patients with Parkinson's disease (Bantis, Nakas, Reiser, Myall, & Dalrymple-Alford 2017). Based on a battery of tests for characterising cognitive impairment, 170 patients were diagnosed as unimpaired (U), 52 patients were diagnosed as having mild cognitive impairment (MCI), and 23 patients were diagnosed as having dementia (D). Parkinson's disease patients who have dementia were expected to have slower completion times than those with MCI, and patients with MCI were expected to have slower completion times than those with no cognitive impairment, that is, the anticipated ordering of completion times is $U < MCI < D$. Figure 5 shows histograms and boxplots of the completion times for each group. We can observe a very reasonable separation between completion times in the three groups, with an almost non-existing overlap between completion times in the unimpaired and dementia group.

We applied our BB methodology to the TMT Part A data. We used 5000 iterations and, as in Section 4, p_1 and p_3 lie on a grid of 50 even points on $[0, 1]$. The BB estimates of the cumulative distribution function in each group along with 95% pointwise credible bands presented in Figure 6 of the Supplementary Materials show, consistent with the histograms and boxplots of the completion times, considerable separation between the distributions of completion times in the unimpaired and mild cognitive impairment groups compared to the severe dementia group. The BB estimate of the ROC surface has the appealing feature of being smooth (without the need for specifying a smoothing parameter), therefore allowing for useful interpretation of diagnostic performance at all threshold values (Figure 6a). Figure 6b presents a histogram of the 5000 sampled VUS values; the BB estimate (95% credible interval) is 0.75 (0.65, 0.83), which indicates that completion time on Part A of the TMT accurately discerns between U, MCI, and D in Parkinson disease patients.

We also applied the kernel and empirical estimators in the same manner as described in Section 4. Confidence intervals for the VUS were obtained through a nonparametric bootstrap consisting of 1000 resamples. The estimated surfaces are shown in Figure 6 (c)–(e). Note that the empirical surface lacks the smoothness property, while the kernel approach achieves it but at the cost of using a bandwidth parameter. The corresponding VUS estimates are 0.70 (0.62, 0.79), 0.67 (0.63, 0.80), and 0.74 (0.66, 0.83) for the kernel method with bandwidth calculated using equation (2), the kernel method with bandwidth selected by cross validation, and the empirical method, respectively. The empirical VUS is similar to our BB estimate, whereas the kernel VUS, for both bandwidths, is slightly lower. This is in agreement to what has been reported by Bantis et al. (2017). Overall, all estimates are similar and suggest that TMT Part A completion time is an accurate test for dementia stage in Parkinson's disease patients.

6 | CONCLUDING REMARKS

We have developed a flexible, nonparametric method based on the Bayesian bootstrap and on the notion of placement value for conducting inference about the ROC surface and its functionals. In addition to providing point and interval estimates in a single integrated framework, our method is computationally easy to implement and very fast. A simulation study illustrated the ability of our approach to produce accurate estimates for a variety of data-generating distributions, and it demonstrated that our estimator is a viable alternative to current nonparametric surface estimators. Furthermore, the validity of our inferences, in terms of frequentist probability of coverage, was demonstrated. The TMT data analysis revealed high accuracy for distinguishing between Parkinson's disease patients who present normal cognition, mild cognitive impairment and dementia. An interesting avenue for future research is the potential use of the Bayesian bootstrap for learning about the ROC surface of tests subject to a limit of detection.

References

- Bantis, L. E., Nakas, C. T., Reiser, B., Myall, D., & Dalrymple-Alford, J. C. (2017). Construction of joint confidence regions for the optimal true class fractions of receiver operating characteristic (ROC) surfaces and manifolds. *Statistical Methods in Medical Research*, 26(3), 1429–1442.
- Branscum, A. J., Johnson, W. O., & Baron, A. T. (2013). Robust medical test evaluation using flexible Bayesian semiparametric regression models. *Epidemiology Research International*, 2013.
- Branscum, A. J., Johnson, W. O., Hanson, T. E., & Baron, A. T. (2015). Flexible regression models for ROC and risk analysis, with or without a gold standard. *Statistics in Medicine*, 34(30), 3997–4015.
- Branscum, A. J., Johnson, W. O., Hanson, T. E., & Gardner, I. A. (2008). Bayesian semiparametric ROC curve estimation and disease diagnosis. *Statistics in Medicine*, 27(13), 2474–2496.
- Coolen-Maturi, E. F., T., & Coolen, F. (2014). Three-group ROC analysis: A nonparametric predictive approach. *Computational Statistics and Data Analysis*, 78, 69–81.
- Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7(1), 1–26.

- Erkanli, A., Sung, M., Jane Costello, E., & Angold, A. (2006). Bayesian semi-parametric ROC analysis. *Statistics in Medicine*, 25(22), 3905–3928.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2), 209–230.
- Gasparini, M. (1995). Exact multivariate Bayesian bootstrap distributions of moments. *The Annals of Statistics*, 23(3), 762–768.
- Gu, J., Ghosal, S., & Roy, A. (2008). Bayesian bootstrap estimation of ROC curve. *Statistics in Medicine*, 27(26), 5407–5420.
- Hong, C., & Cho, M. (2015). Test statistics for volume under the ROC surface and hypervolume under the ROC manifold. *Communications for Statistical Applications and Methods*, 22, 377–387.
- Hwang, B. S., & Chen, Z. (2015). An integrated Bayesian nonparametric approach for stochastic and variability orders in ROC curve estimation: an application to endometriosis diagnosis. *Journal of the American Statistical Association*, 110(511), 923–934.
- Inácio, V., Turkman, A. A., Nakas, C. T., & Alonzo, T. A. (2011). Nonparametric Bayesian estimation of the three-way receiver operating characteristic surface. *Biometrical Journal*, 53(6), 1011–1024.
- Inácio de Carvalho, V., & Branscum, A. J. (2018). Bayesian nonparametric inference for the three-class Youden index and its associated optimal cutoff points. *Statistical Methods in Medical Research*, 27(3), 689–700.
- Inácio de Carvalho, V., de Carvalho, M., & Branscum, A. J. (2017). Nonparametric Bayesian covariate-adjusted estimation of the Youden index. *Biometrics*, 73(4), 1279–1288.
- Inácio de Carvalho, V., Jara, A., Hanson, T. E., & de Carvalho, M. (2013). Bayesian nonparametric ROC regression modeling. *Bayesian Analysis*, 8(3), 623–646.
- Kang, L., & Tian, L. (2013). Estimation of the volume under the ROC surface with three ordinal diagnostic categories. *Computational Statistics & Data Analysis*, 62, 39–51.
- Kim, Y., Lee, J., & Kim, J. (2005). Bayesian bootstrap analysis of doubly censored data using Gibbs sampler. *Statistica Sinica*, 969–980.
- Li, J., Zhou, X., & Fine, J. P. (2012). A regression approach to ROC surface, with applications to Alzheimer's disease. *Science China Mathematics*, 55, 1583–1595.
- Li, J., & Zhou, X.-H. (2009). Nonparametric and semiparametric estimation of the three way receiver operating characteristic surface. *Journal of Statistical Planning and Inference*, 139(12), 4133–4142.
- Nakas, C. T. (2014). Developments in ROC surface analysis and assessment of diagnostic markers in three-class classification problems. *REVSTAT*, 12, 43–65.
- Nakas, C. T., & Yiannoutsos, C. T. (2004). Ordered multiple-class ROC analysis with continuous measurements. *Statistics in Medicine*, 23, 3437–3449.
- Pepe, M. S. (2003). *The Statistical Evaluation of Medical Tests for Classification and Prediction*. Oxford University Press, New York.
- R Core Team. (2017). *R: A language and environment for statistical computing* [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Rodríguez, A., & Martínez, J. C. (2013). Bayesian semiparametric estimation of covariate-dependent ROC curves. *Biostatistics*, 15(2), 353–369.
- Rubin, D. B. (1981). The Bayesian bootstrap. *The Annals of Statistics*, 9(1), 130–134.
- Wan, S. (2012). An empirical likelihood confidence interval for the volume under ROC surface. *Statistics & Probability Letters*, 82, 1463–1467.
- Wand, M. P., & Jones, M. C. (1994). *Kernel smoothing*. CRC/Chapman Hall, Boca Raton.
- Yang, H., & Carlin, D. (2000). ROC surface: a generalization of ROC curve analysis. *Journal of Biopharmaceutical Statistics*, 10, 183–196.
- Yu, T. (2012). ROCs: Receiver operating characteristic surface for class-skewed high-throughput data. *PLOS One*, 7(7).
- Zhang, Y., & Alonzo, T. (2016). Inverse probability weighting estimation of the volume under the ROC surface in the presence of verification bias. *Biometrical Journal*, 58, 1338–1356.
- Zhang, Y., & Li, J. (2011). Combining multiple markers for multi-category classification: An ROC surface approach. *Australian & New Zealand Journal of Statistics*, 53, 63–78.
- Zhao, L., Feng, D., Chen, G., & Taylor, J. M. (2016). A unified Bayesian semiparametric approach to assess discrimination ability in survival analysis. *Biometrics*, 72(2), 554–562.
- Zhou, X.-H., McClish, D. K., & Obuchowski, N. A. (2011). *Statistical methods in diagnostic medicine* (2nd ed.). John Wiley & Sons, New York.

Scenario	Y_1	Y_2	Y_3
1	$N(0, 1^2)$	$N(1.5, 1^2)$	$N(3, 1^2)$
2	$\text{Gamma}(2, 1)$	$\text{Gamma}(3, 1)$	$\text{Gamma}(5, 2)$
3	t_2	$\text{Beta}(2, 2)$	χ_2^2
4	$0.5N(0, 1^2) + 0.5N(3, 1^2)$	$0.5N(1, 1^2) + 0.5N(4, 1.5^2)$	$0.5N(2, 1^2) + 0.5N(5, 2^2)$

TABLE 1 Scenarios considered for the simulation study.

Scenario	$(n_1, n_2, n_3) = (50, 50, 50)$	$(n_1, n_2, n_3) = (100, 100, 100)$	$(n_1, n_2, n_3) = (200, 200, 200)$
1	0.95	0.98	0.97
2	0.99	1	0.99
3	0.95	0.97	0.95
4	0.98	0.98	0.98

TABLE 2 VUS 95% coverage probabilities.

FIGURE 1 Scenario 1. Boxplots summarising simulation results for the EMSE (top row) and estimates of the VUS (bottom row). The solid red line corresponds to the true VUS. Here Kernel denotes the kernel estimate with bandwidth calculated using equation (2) and Kernel-CV stands for the kernel estimate with the bandwidth selected by least squares cross-validation.

FIGURE 2 Scenario 2. Boxplots summarising simulation results for the EMSE (top row) and estimates of the VUS (bottom row). The solid red line corresponds to the true VUS. Here Kernel denotes the kernel estimate with bandwidth calculated using equation (2) and Kernel-CV stands for the kernel estimate with the bandwidth selected by least squares cross-validation.

FIGURE 3 Scenario 3. Boxplots summarising simulation results for the EMSE (top row) and estimates of the VUS (bottom row). The solid red line corresponds to the true VUS. Here Kernel denotes the kernel estimate with bandwidth calculated using equation (2) and Kernel-CV stands for the kernel estimate with the bandwidth selected by least squares cross-validation.

FIGURE 4 Scenario 4. Boxplots summarising simulation results for the EMSE (top row) and estimates of the VUS (bottom row). The solid red line corresponds to the true VUS. Here Kernel denotes the kernel estimate with bandwidth calculated using equation (2) and Kernel-CV stands for the kernel estimate with the bandwidth selected by least squares cross-validation.

FIGURE 5 Trail Making Test Part A data: histograms and variable-width boxplots of the completion times (in seconds) in each group.

FIGURE 6 Trail Making Test Part A data. Top row: BB estimate of the ROC surface (a) and histogram of the 5000 BB sampled VUS (b). Bottom row: estimated ROC surfaces using the kernel approach with bandwidths computed as in (2) (c), kernel approach with bandwidths selected by cross-validation (d), and the empirical approach (e).