

Model-free posterior inference on the area under the receiver operating characteristic curve

Zhe Wang* and Ryan Martin*

August 2, 2020

Abstract

The area under the receiver operating characteristic curve (AUC) serves as a summary of a binary classifier's performance. For inference on the AUC, a common modeling assumption is *binormality*, which restricts the distribution of the score produced by the classifier. However, this assumption introduces an infinite-dimensional nuisance parameter and may be restrictive in certain machine learning settings. To avoid making distributional assumptions, and to avoid the computational challenges of a fully nonparametric analysis, we develop a direct and model-free Gibbs posterior distribution for inference on the AUC. We present the asymptotic Gibbs posterior concentration rate, and a strategy for tuning the learning rate so that the corresponding credible intervals achieve the nominal frequentist coverage probability. Simulation experiments and a real data analysis demonstrate the Gibbs posterior's strong performance compared to existing Bayesian methods.

Keywords and phrases: credible interval; Gibbs posterior; generalized Bayesian inference; model misspecification; robustness.

1 Introduction

First proposed during World War II to assess the performance of radar receiver operators (Calì and Longobardi 2015), the receiver operating characteristic (ROC) curve is now an essential tool for analyzing the performance of binary classifiers in areas such as signal detection (Green and Swets 1966), psychology examination (Swets 1973, 1986), radiology (Hanley and McNeil 1982; Lusted 1960), medical diagnosis (Hanley 1989; Swets and Pickett 1982), and data mining (Fawcett 2006; Spackman 1989). One informative summary of the ROC curve is the corresponding area under the curve (AUC). This measure provides an overall assessment of classifier's performance, independent of the choice of threshold, and is, therefore, the preferred method for evaluating classification algorithms (Bradley 1997; Huang and Ling 2005; Provost and Fawcett 1997; Provost et al. 1998). The AUC is an unknown quantity, and our goal is to use the information contained in the data to make inference about the AUC. The specific set up is as follows. For a binary classifier which produces a random *score* to indicate the propensity for, say,

*Department of Statistics, North Carolina State University; zwang54@ncsu.edu, rgmarti3@ncsu.edu

Group 1; individuals with scores higher than a threshold are classified to Group 1, the rest are classified to Group 0. Let U and V be independent scores corresponding to Group 1 and Group 0, respectively. Given a threshold t , define the *specificity* and *sensitivity* as $\text{spec}(t) = \mathbb{P}(V < t)$ and $\text{sens}(t) = \mathbb{P}(U > t)$. Then the ROC curve is a plot of the parametric curve $(1 - \text{spec}(t), \text{sens}(t))$ as t takes all possible values for scores. While the ROC curve summarizes the classifier's tradeoff between sensitivity and specificity as the threshold varies, the AUC measures the probability of correctly assigning scores for two individuals from two groups, which equals $\mathbb{P}(U > V)$ (Bamber 1975), and is independent of the choice of threshold. Consequently, the AUC is a functional of the joint distribution of (U, V) , denoted by \mathbb{P} , so the ROC curve is actually not needed to identify AUC.

In the context of inference on the AUC, when the scores are continuous, it is common to assume that \mathbb{P} satisfies a so-called *binormality* assumption, which states that there exists a monotone increasing transformation that maps both U and V to normal random variables (Hanley 1988). For most medical diagnostic tests, where the classifiers are simple and ready-to-use without training, such an assumption serves well (Cai and Moskowitz 2004; Hanley 1988; Metz et al. 1998), although it has been argued that other distributions can be more appropriate for some specific tests (e.g., Goddard and Hinberg 1990; Guignard and Salehi 1983). But for complicated classifiers which involve multiple predictors, as often arise in machine learning applications, binormality—or any other model assumption for that matter—becomes a burden. This motivates our pursuit of a “model-free” approach to inference about the AUC.

Specifically, our goal is the construction of a type of posterior distribution for the AUC. The most familiar such construction is via Bayes's formula, but this requires a likelihood function and, hence, a statistical model. The only way one can be effectively “model-free” within a Bayesian framework is to make the model extra flexible, which requires lots of parameters. In the extreme case, a so-called Bayesian nonparametric approach would take the distribution \mathbb{P} itself as the model parameter (e.g., Ghosal and van der Vaart 2017; Gu et al. 2008). When the model includes lots of parameters, then the analyst has the burden of specifying prior distributions for these, based on little or no genuine prior information, and also computation of a high-dimensional posterior. But since the AUC is just a one-dimensional feature of this complicated set of parameters, there is no obvious return on the investment into prior specification and posterior computation. A better approach would be to construct the posterior distribution for the AUC directly, using available prior information about the AUC only, without specifying a model and without the introduction of artificial model parameters. That way, the data analyst can avoid the burdens of prior specification and posterior computation, bias due to model misspecification, and issues that can arise as a result of non-linear marginalization (e.g., Fraser 2011; Martin 2019).

As an alternative to the traditional Bayesian approach, we consider here the construction of a so-called *Gibbs posterior* for the AUC. In general, the Gibbs posterior construction proceeds by defining the quantity of interest as the minimizer of a suitable risk function, treating an empirical version of that loss function like a negative log-likelihood, and then combining with a prior distribution like in Bayes's formula. General discussion of Gibbs posteriors can be found in Zhang (2006a,b), Bissiri et al. (2016) and Alquier et al. (2016), and some statistical applications are discussed in Jiang and Tanner (2008) and Syring and Martin (2017, 2019a,b). Again, the advantage is that Gibbs posteriors avoid

model misspecification bias and the need to deal with nuisance parameters. Moreover, under suitable conditions, Gibbs posteriors can be shown to have desirable asymptotic concentration properties (e.g., Bhattacharya and Martin 2020; Chernozhukov and Hong 2003; Syring and Martin 2020), with theory that parallels that of Bayesian posteriors under model misspecification (e.g., Kleijn and van der Vaart 2006, 2012).

A subtle point is that, while the risk minimization problem that defines the quantity of interest is independent of the scale of the loss function, the Gibbs posterior is not. This scale factor is often referred to as the *learning rate* (e.g., Grünwald 2012) and, because it controls the spread of the Gibbs posterior, its specification needs to be handled carefully. Various approaches to the specification of the learning rate parameter (e.g., Bissiri et al. 2016; Grünwald 2012; Grünwald and Van Ommen 2017; Holmes and Walker 2017; Lyddon et al. 2019). Here we adopt the approach in Syring and Martin (2019a) that aims to set the learning rate so that, in addition to its robustness to model misspecification and asymptotic concentration properties, the Gibbs posterior credible sets have the nominal frequentist coverage probability. When the sample size is large, we recommend an (asymptotically) equivalent calibration method that is simpler to compute.

The present paper is organized as follows. In Section 2.1, we review some methods for making inference on the AUC based on the binormality assumption, in particular, the Bayesian approach in Gu and Ghosal (2009) that involves a suitable rank-based likelihood. In Section 2.2, we argue that the binormality assumption is generally inappropriate in machine learning applications, and provide one illustrative example involving a support vector machine. This difficulty with model specification leads us to the Gibbs posterior, a model-free alternative to a Bayesian posterior, which is reviewed in Section 2.3. We develop the Gibbs posterior for inference on the AUC, derive its asymptotic concentration properties, and investigate how to properly scale the risk function in Section 3. Simulation experiments are carried out in Section 4, where a Gibbs posterior estimator performs favorably compared with the Bayesian approach based on a rank-based likelihood. We also apply the Gibbs posterior on a real dataset for evaluating the performance of a biomarker for pancreatic cancer and compare our result with those based on the rank likelihood. Finally, we give some concluding remarks in Section 5.

2 Background

2.1 Binormality and related methods

Following Hanley (1988), the scores U and V satisfy the binormality assumption if their distribution functions are $\Phi[b^{-1}\{H(u) - a\}]$ and $\Phi\{H(v)\}$ respectively, where $a > 0$, $b > 0$, H is a monotone increasing function, and Φ denotes the $\mathcal{N}(0, 1)$ distribution function, which implies that U and V can be transformed to $\mathcal{N}(a, b^2)$ and $\mathcal{N}(0, 1)$ via H . If $\mathbb{P} = \mathbb{P}_{a,b,H}$ denotes the distribution of (U, V) under this assumption, then the ROC curve and the AUC, respectively, are given by $t \mapsto \Phi[b^{-1}\{a + \Phi^{-1}(t)\}]$ and

$$\Phi\{a(b^2 + 1)^{-1/2}\}. \quad (1)$$

Even though H is not needed to define the AUC—only (a, b) —since the joint distribution of (U, V) does depend on H , any likelihood-based method would have to deal with this

infinite-dimensional nuisance parameter. Some strategies are used to avoid dealing with H directly. The semi-parametric approach in Cai and Moskowitz (2004) manipulates the equivalent densities ratio of U over V and W over Z , and introduces cumulative hazard function as a nuisance parameter. A profile likelihood is obtained based on a discrete estimate for the cumulative hazard function. In the approach of Metz et al. (1998), data are suitably grouped and a multinomial pseudo-likelihood is constructed. Alternatively, since data ranks are invariant to monotone transformations, one can construct a rank-based likelihood, as in Zou and Hall (2000), which can be maximized over (a, b) to estimate the AUC. But it turns out that a Bayesian approach that uses Monte Carlo sampling from a rank-based posterior distribution, as in Gu and Ghosal (2009), is computationally more efficient than maximizing the rank likelihood. Since this is our proposed method's primary competitor, we give some details about Gu and Ghosal's Bayesian rank-based likelihood approach here.

Consider the transformed scores $W = H(U)$ and $Z = H(V)$, according to the binormality assumption, its joint distribution can be written as $\mathbb{P}_{a,b}$, no more dependence on H . Elimination of the nuisance parameter H is desirable, but (W, Z) are unavailable to us without knowledge of H . That is, unless we consider a function of (U, V) that is invariant to transformations by H . A good candidate function is the *ranks*. That is, let $R_{U,V}$ denote the ranks of the vector $(U_1, \dots, U_m, V_1, \dots, V_n)$, where (U_1, \dots, U_m) and (V_1, \dots, V_n) are independent and identically distributed (iid) copies of U and V , respectively. Then

$$\mathbb{P}_{a,b,H}(R_{U,V} = r) \equiv \mathbb{P}_{a,b}(R_{W,Z} = r), \quad (2)$$

where $R_{W,Z}$ is the ranks of $(W_1, \dots, W_m, Z_1, \dots, Z_n)$, with $W_i = H(U_i)$ and $Z_j = H(V_j)$. The key is that the observed ranks based on the (U_i, V_j) sample can be plugged in for r on the right-hand side of (2) and that gives a likelihood function for (a, b) , without requiring knowledge of H . Of course, this is not a proper likelihood function, i.e., there is loss of information caused by throwing away the values of (U_i, V_j) , but eliminating the infinite-dimensional nuisance parameter might be worth the price, especially when the goal is inference on the ROC curve or AUC, neither of which depend directly on H . The approach outlined in Gu and Ghosal (2009) proceeds by treating the (W_i, Z_j) values as latent variables and defining a full posterior for (a, b^2, W_i, Z_j) , given $R_{U,V}$, and then marginalizing out (W_i, Z_j) to get a posterior distribution for (a, b^2) alone. If we take the Jeffreys prior for (a, b^2) , which is proportional to b^{-2} , then the full conditional distribution presented in Gu and Ghosal (2009) are

$$\begin{aligned} (a \mid W_1, \dots, W_m, Z_1, \dots, Z_n, b^2, R_{U,V}) &\sim \mathcal{N}(m^{-1} \sum_{i=1}^m W_i, b^2 m^{-1}), \\ (b^2 \mid W_1, \dots, W_m, Z_1, \dots, Z_n, a, R_{U,V}) &\sim \text{IG}\left(\frac{m-1}{2}, \frac{1}{2} \sum_{i=1}^m (W_i - a)^2\right), \\ (W_i \mid W_{-i}, Z_1, \dots, Z_n, a, b^2, R_{U,V}) &\sim \mathcal{N}(a, b^2) \times 1(R_{W,Z} = R_{U,V}), \quad i = 1 \dots m \\ (Z_j \mid Z_{-j}, W_1, \dots, W_m, a, b^2, R_{U,V}) &\sim \mathcal{N}(0, 1) \times 1(R_{W,Z} = R_{U,V}), \quad j = 1 \dots n \end{aligned}$$

where, e.g., $W_{-i} = (W_1, \dots, W_{i-1}, W_{i+1}, \dots, W_m)$, $\text{IG}(\alpha, \beta)$ denotes the inverse gamma distribution with density $\frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} e^{-\beta/x}$, and $1(\cdot)$ denotes the indicator function. With these full conditionals, it is straightforward to develop a Monte Carlo strategy that produces samples from the (a, b^2) posterior distribution. These samples can then be used to get a posterior distribution for AUC using the expression in (1).

2.2 Validity of binormality in machine learning applications

Before ROC and AUC analyses gained popularity machine learning, the binormality assumption had been proposed and used in the context of medical diagnosis for *simple classifiers*, where the scores U and V are determined based on a single predictor variable. When assuming binormality for classifiers in machine learning (e.g., Brodersen et al. 2010; Macskassy and Provost 2004), the situation differs because multiple predictor variables are usually involved.

Suppose that a binary $Y \in \{0, 1\}$ indicates the group, $X \in \mathbb{R}^p$ is the predictors, and $\mathcal{M}_{X,Y}$ denotes the joint distribution of (X, Y) . As described in Section 1, a binary classifier provides a parametric form of the predictor, namely the *score* $S(X; \beta)$, to indicate the propensity for Y taking value 1. A training process is generally needed for estimating the unknown β based on data $\{(X_i, Y_i) : i = 1, \dots, n\}$, and let $\hat{\beta}_n \in \mathbb{R}^p$ denote that estimator. Then the random score U for Group 1 is defined as $S(X; \hat{\beta}_n)$ where the predictor X follows the conditioned distribution $\mathcal{M}_{X|Y=1}$. Similarly, the random score V for Group 0 is defined as $S(X; \hat{\beta}_n)$ where $X \sim \mathcal{M}_{X|Y=0}$. By assuming that $\hat{\beta}_n$ converges to a non-random quantity as $n \rightarrow \infty$, U and V are asymptotically independent. Consequently, for simple classifiers, where $p = 1$ and $S(X; \beta) = X$, a binormality assumption on U and V only restricts $\mathcal{M}_{X|Y=1}$ and $\mathcal{M}_{X|Y=0}$ to be two normal distributions (after the transformation H). In the general case where $p > 1$, even if each predictor is normally distributed, because $S(X; \beta)$ might be a complicated non-linear function, the corresponding scores U and V need not satisfy binormality.

For example, we suppose that predictors X_1 and X_2 are used to classify two groups ($Y = 1$ or 0). In Group 1, X_1 and X_2 are independently and identically distributed as $\mathcal{N}(0, 50)$; in Group 0, they are distributed as $\mathcal{N}(0, 2)$. Panel (a) in Figure 1 shows the training data, and it is clear that a non-linear classifier would be required to separate the two groups. Here we use a support vector machine with the radial basis function kernel; based on the training result, we obtain scores $\{U_i\}_{i=1}^m$ and $\{V_j\}_{j=1}^n$ for each group, $m = n = 10,000$. To verify the binormality, we approximate the monotone increasing transformation H by $H = \Phi^{-1} \circ \hat{F}_V$, where $\hat{F}_V(v) = n^{-1} \sum_{j=1}^n 1(V_j \leq v)$ is the empirical distribution function. Although H uniquely transforms V to a standard normal, the transformed score $H(U)$ does not follow a normal distribution, as indicated in Panels (b) and (c) of Figure 1. Therefore, given this specific dataset and classifiers, there is no such a H which simultaneously transforms U and V to normal random variables.

2.3 Gibbs posterior distributions

A Gibbs posterior distribution resembles a Bayesian posterior, but is constructed using different ingredients. In particular, the Gibbs posterior does not start with a statistical model and likelihood, it starts with a more general connection between data and quantities of interest, through a loss function. Suppose that data T_1, \dots, T_n are identically distributed \mathbb{T} -valued observations from distribution P , and that there is some functional $\theta = \theta(\mathsf{P})$, taking values in Θ , about which inference is desired. Instead of introducing a statistical model for P —that is, assuming P takes a particular distributional form P_ζ for some model parameter ζ , and then expressing θ as a function of ζ —we construct a posterior for θ directly as follows. Assume that there exists a loss function $\ell_\theta(t)$, mapping

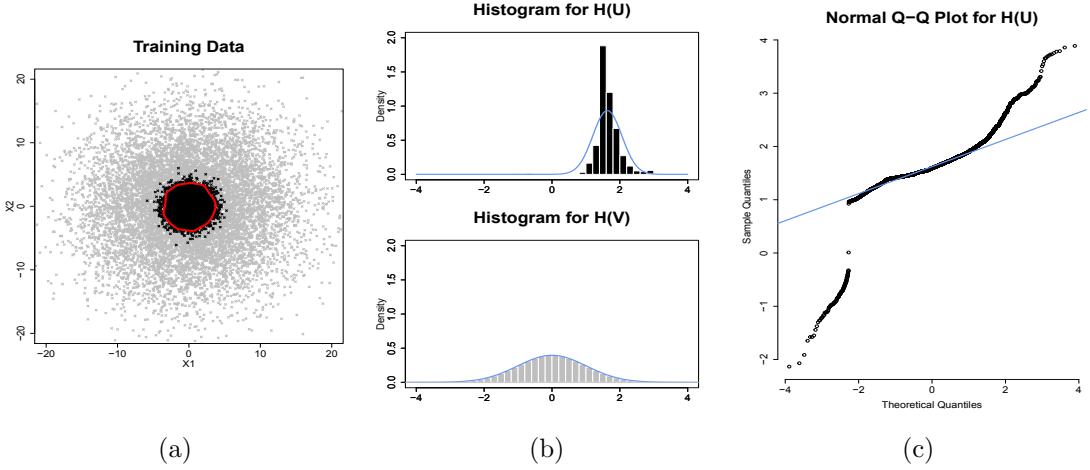


Figure 1: (a) Training set and the SVM decision boundary (red curve). For each point, two predictors (X_1, X_2) are plotted on axes and Y is visualized by the color (black for Group 1, gray for Group 0); (b) Histograms for transformed scores $H(U)$ and $H(V)$ with the fitted normal densities; (c) Q-Q plot for transformed samples $\{H(U_1), \dots, H(U_m)\}$

$\mathbb{T} \times \Theta$ to \mathbb{R} , such that the true value, θ^* , of θ solves the optimization problem

$$\theta^* = \arg \min_{\theta} R(\theta), \quad (3)$$

where the risk function $R(\theta) = \mathbb{P}\ell_{\theta}$ is just the expected loss with respect to \mathbb{P} . When the quantity of interest is defined as the solution to an optimization problem, it makes sense to estimate that quantity by solving an empirical version of the optimization problem,

$$\hat{\theta}_n = \arg \min_{\theta} R_n(\theta),$$

where the empirical risk $R_n(\theta) = \hat{\mathbb{P}}_n \ell_{\theta}$ is the expected loss with respect to the empirical distribution $\hat{\mathbb{P}}_n = n^{-1} \sum_{i=1}^n \delta_{T_i}$, with δ_t the point-mass distribution concentrated at t . From this empirical risk function, the Gibbs posterior distribution is defined as

$$\Pi_n(d\theta) \propto e^{-\omega n R_n(\theta)} \Pi(d\theta), \quad \theta \in \Theta, \quad (4)$$

where Π is a prior distribution on Θ and $\omega > 0$ is a scale parameter to be determined; see Bissiri et al. (2016) for the decision-theoretic underpinnings of this approach.

For us, the motivation behind the use of a Gibbs posterior is that it gives us direct, model-free posterior inference about the quantity of interest. This is beneficial because, for one thing, a statistical model could be misspecified and that would generally bias the results. But even if the model is correctly specified, it is unlikely that an appropriate statistical model could be described in terms of θ alone, so the model index ζ would include a number of nuisance parameters that require prior distribution specification and posterior computation, efforts that are effectively wasted if marginal inference on θ is the goal. The Gibbs posterior, by targeting θ directly, avoids the possible misspecification bias, allows for prior beliefs about θ to be readily accommodated, and does not require dealing with nuisance parameters. And the applications presented in Syring and Martin

(2017, 2019a,b), along with the one presented here, suggest that this direct approach has a number of important advantages over the more traditional Bayesian counterpart.

Of course, the magnitude of the loss function does not affect the solution to the optimization problem in (3), nor that in the empirical version thereof. But the magnitude does affect the Gibbs posterior in (4), which is why we include the scaling factor ω . Data-driven strategies for specifying this tuning parameter are discussed in Section 3.3 below.

3 Gibbs posterior for the AUC

3.1 Definition

As mentioned, the AUC is a functional of the joint distribution P of (U, V) , i.e., $\theta = \theta(\mathsf{P})$, given by $\theta = \mathsf{P}(U > V)$. Recall that the data consists of independent copies (U_1, \dots, U_m) and (V_1, \dots, V_n) of U and V , respectively. To construct a Gibbs posterior distribution for θ as discussed above, we need an appropriate loss function. That is, we need a function $\ell_\theta(u, v)$ such that the corresponding risk function, $R(\theta) = \mathsf{P}\ell_\theta$, is minimized at the true AUC, θ^* . If we define

$$\ell_\theta(u, v) = \{\theta - 1(u > v)\}^2, \quad \theta \in [0, 1],$$

then it is easy to check that $R(\theta) = \theta^2 - 2\theta^*\theta + \theta^{*2}$ and, moreover, that this risk function is uniquely minimized at $\theta = \theta^*$. Then the empirical risk function is

$$R_{m,n}(\theta) = \widehat{\mathsf{P}}_{m,n}\ell_\theta = \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n \{\theta - 1(U_i > V_j)\}^2$$

where $\widehat{\mathsf{P}}_{m,n} = (mn)^{-1} \sum_{i=1}^m \sum_{j=1}^n \delta_{(U_i, V_j)}$ is the empirical distribution of the score pairs. Note that the minimizer of the empirical risk function, namely,

$$\hat{\theta}_{m,n} = \arg \min_{\theta} R_{m,n}(\theta) = \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n 1(U_i > V_j), \quad (5)$$

is the familiar statistic suggested by Mann and Whitney (1947) for testing if one of two independent random variables is stochastically larger than the other.

Following the general approach described in Section 2.3, we can construct a Gibbs posterior distribution for the AUC, with density

$$\pi_{m,n}(\theta) \propto e^{-\omega mn R_{m,n}(\theta)} \pi(\theta), \quad \theta \in [0, 1],$$

where π is some prior density for the AUC, and ω is the learning rate to be specified in Section 3.3. This Gibbs posterior does not require any model assumptions, does not require marginalization over nuisance parameters, and can directly incorporate available prior information about θ . Moreover, the Gibbs posterior is approximately centered around $\hat{\theta}_{m,n}$, which is a quality estimator of the AUC, regardless of what form the underlying distribution P takes, so we can expect the Gibbs posterior—for suitable ω —to provide quality model-free inference. Some nonstandard concentration results emerge, see Theorem 1, because the summation in $R_{m,n}$ contains $m \times n$ dependent terms. To

construct the Gibbs posterior, the empirical risk function is multiplied by the product of two sample sizes $m \times n$ rather than its total $m + n$.

After some simple algebraic manipulations, the Gibbs posterior above can be re-expressed as

$$\pi_{m,n}(\theta) \propto e^{-\omega mn(\theta - \hat{\theta}_{m,n})^2} \pi(\theta), \quad \theta \in [0, 1], \quad (6)$$

which shows some resemblance to a truncated normal distribution. A very reasonable choice of prior is a truncated normal distribution with informative choices of prior location μ_0 and scale σ_0 . With this choice, the Gibbs posterior is a truncated normal distribution too, with corresponding location and scale, respectively,

$$\mu_{m,n} = \frac{\mu_0 + 2\omega\sigma_0^2 mn \hat{\theta}_{m,n}}{1 + 2\omega\sigma_0^2 mn} \quad \text{and} \quad \sigma_{m,n} = \left\{ \frac{\sigma_0^2}{1 + 2\omega\sigma_0^2 mn} \right\}^{1/2}.$$

In the absence of prior information about the AUC, one can take a flat uniform prior, $\pi(\theta) \equiv 1$, in which case the Gibbs posterior is still a truncated normal distribution but with location and scale, respectively,

$$\mu_{m,n} = \hat{\theta}_{m,n} \quad \text{and} \quad \sigma_{m,n} = (2\omega mn)^{-1/2}.$$

In practice, we recommend the use of available prior information about the AUC whenever possible, but, for the remainder of this paper, we will work with the Gibbs posterior based on the default uniform prior.

Here we are interested in the inference on AUC for any given classifier, and consequently the posterior is constructed directly for the AUC. A similar Gibbs posterior formulation can be found in Ridgway et al. (2014), which aimed at selecting the classifier via AUC maximization and focused on the parameters introduced by the classifiers.

3.2 Asymptotic concentration properties

It is natural to ask what kind of asymptotic concentration properties the Gibbs posterior distribution enjoys. An advantage of our approach's simplicity is the ease in which the convergence properties can be deduced, but some care is needed in formulating the asymptotic regime precisely. Indeed, since the two groups may have different sample sizes, it is clear that what we need is for the *smaller* of the two sample sizes to go to infinity. Therefore, the rate is determined by $m \wedge n$, and following theorem states that, under no conditions on the joint distribution \mathbf{P} of (U, V) , the Gibbs posterior distribution concentrates asymptotically around the true AUC at the rate $(m \wedge n)^{-1/2}$.

Theorem 1. *Let θ^* be the true AUC corresponding to \mathbf{P} . If $\Pi_{m,n}$ is the Gibbs posterior defined in (6) based on a fixed learning rate $\omega > 0$ and a prior density π that is positive and continuous in an interval containing θ^* , then for any sequence $K_{m,n} \rightarrow \infty$,*

$$\Pi_{m,n}(\{\theta : |\theta - \theta^*| > K_{m,n} (m \wedge n)^{-1/2}\}) \rightarrow 0 \quad \text{in } \mathbf{P}\text{-probability as } m \wedge n \rightarrow \infty.$$

Proof. See Appendix A. □

Several remarks on the concentration rate theorem, its consequences, and some related results are in order.

- The convergence in P -probability conclusion in Theorem 1 implies convergence in $L_1(\mathsf{P})$ by the dominated convergence theorem, and can be strengthened to convergence with P -probability 1 by assuming that sample sizes for two groups increase at the same rate, i.e., $m(m+n)^{-1} \rightarrow \rho \in (0, 1)$. Under this condition, Korolyuk and Borovskich (2013, Chap. 3.2) show that $\hat{\theta}_{m,n} \rightarrow \theta^*$ with P -probability 1 and, with this, the stronger Gibbs posterior concentration rate result can be proved along lines similar to those in Appendix A below.
- As shown in (4), the Gibbs posterior resembles a Bayesian posterior based on a suitably misspecified model, one whose “likelihood function” equals $\exp\{-\omega n R_n(\theta)\}$. Even in misspecified cases, Bernstein–von Mises-style distributional approximations are possible; see, e.g., Kleijn and van der Vaart (2012). In our case, we immediately see a truncated normal form of the Gibbs posterior, so as long as $\theta^* \in (0, 1)$, the asymptotic normality of the Gibbs posterior is automatic.
- We note the loss scale ω controls the proportion of information in the Gibbs posterior which is learned from the data. Consequently, it is reasonable to adjust ω so that a set of observations with a larger size is given more trust. In fact, if we substitute the fixed ω in Theorem 1 with a sequence ω_n that vanishes slower than $(m \vee n)^{-1}$, then the Gibbs posterior concentration rate result still holds.

3.3 Tuning the learning rate

The good behavior of a Bayesian posterior is guaranteed only when the model is correctly specified. Under misspecification, even if the posterior concentrate around an efficient estimator, the asymptotic variance of the posterior could be drastically different from that of the efficient estimator; see Kleijn and van der Vaart (2012). Consequently, $100(1 - \alpha)\%$ credible regions from a misspecified Bayes model may not achieve the nominal $100(1 - \alpha)\%$ confidence. Fortunately, the Gibbs posterior learning rate parameter, ω , which controls the spread, can be tuned in such a way that this undesirable discrepancy between credibility and confidence is avoided. Various tuning strategies are available in the literature (e.g., Bissiri et al. 2016; Fasiolo et al. 2017; Grünwald 2012; Lyddon et al. 2019), but only the approach presented in Syring and Martin (2019a) focuses directly on coverage probability, so that is the approach we will adopt here.

Algorithm 1 describes the calibrating procedure from Syring and Martin (2019a) in the context of inference on the AUC. The rationale behind this algorithm is as follows. Take a $100(1 - \alpha)\%$ credible interval based on the Gibbs posterior (6) with learning rate ω , in particular, the highest posterior density credible interval. Then the frequentist coverage probability of that credible interval, call it $c_\alpha(\omega)$, depends on ω , α , and other things. If we could evaluate $c_\alpha(\omega)$, that is, if we knew and could directly simulate from P , then we could just solve the equation $c_\alpha(\omega) = 1 - \alpha$. For future reference, in this ideal case, we call the solution to this equation the *oracle learning rate*. In real applications, however, P is unknown, so we cannot evaluate $c_\alpha(\omega)$ exactly, but we can get an estimate using the bootstrap, and then solve that equation using stochastic approximation (Robbins and Monro 1951) with step size sequence (κ_t) that satisfies

$$\sum_{t=1}^{\infty} \kappa_t = \infty \quad \text{and} \quad \sum_{t=1}^{\infty} \kappa_t^2 < \infty \tag{7}$$

Details are discussed in Syring and Martin (2019a).

When implementing Algorithm 1, for each bootstrap sample and a fixed scale ω , calculating the Gibbs posterior credible interval is straightforward thanks to the simple truncated normal distribution form. However, that credible set calculation is required for every bootstrap sample in each ω updating step, so this can be computationally expensive, especially when the sample is large. Fortunately, for such cases, there is an alternative strategy, which attempts to match the Gibbs posterior variance to the variance of its mode, the Mann–Whitney estimator $\hat{\theta}_{m,n}$. A similar variance-matching idea can be seen in Lyddon et al. (2019). Specifically, if

$$\lambda = \lim_{m,n \rightarrow \infty} \frac{m}{m+n} \in (0, 1),$$

then Hoeffding (1948, Theorem 7.3) showed that the asymptotic variance of $\hat{\theta}_{m,n}$ is

$$\frac{1}{m+n} \left(\frac{\tau_{10}}{\lambda} + \frac{\tau_{01}}{1-\lambda} \right), \quad (8)$$

where

$$\tau_{10} = \mathbf{C}\{1(U_1 > V_1), 1(U_1 > V_2)\} \quad \text{and} \quad \tau_{01} = \mathbf{C}\{1(U_1 > V_1), 1(U_2 > V_1)\},$$

with \mathbf{C} the covariance operator under joint distribution \mathbf{P} . If we take the flat prior in our Gibbs posterior construction, then choosing

$$\hat{\omega}_n = \frac{m+n}{2mn} \left(\frac{\hat{\tau}_{10}}{\lambda} + \frac{\hat{\tau}_{01}}{1-\lambda} \right)^{-1}, \quad (9)$$

with the obvious estimates

$$\begin{aligned} \hat{\tau}_{10} &= \frac{2}{mn(n-1)} \sum_{i=1}^m \sum_{j \neq j'} 1(U_i > V_j) 1(U_i > V_{j'}) - \hat{\theta}_{m,n}^2 \\ \hat{\tau}_{01} &= \frac{2}{nm(m-1)} \sum_{j=1}^n \sum_{i \neq i'} 1(U_i > V_j) 1(U_{i'} > V_j) - \hat{\theta}_{m,n}^2, \end{aligned}$$

will make the Gibbs posterior variance approximately match the Mann–Whitney estimator variance, thus, approximate calibration.

Compared to Algorithm 1, this calibration strategy is computationally efficient; it can also provide an initial guess for Algorithm 1. According to our simulation results in the next section, the two calibration methods behave similarly as the sample size increases, but the optimal scale $\hat{\omega}_n$ yield by (9) has more variability when the sample size is small. Therefore, when the sample size is relatively small, we suggest the strategy in Algorithm 1; on the other hand, when the sample size is large, the other strategy described above is preferable since it achieves similar performance with lower computation cost. The numerical results in Section 4 are all based on the Algorithm 1, and a summary of the ω values selected by the two methods, justifying these claims, is shown in Figure 3.

Algorithm 1: Gibbs posterior calibration

Data: U_1, \dots, U_m and V_1, \dots, V_n

Input: Prior distribution; estimate $\hat{\theta}_{m,n}$ from (5); bootstrap sample size B ; tolerance $\varepsilon > 0$; and step sizes (κ_t) satisfying (7).

Output: An estimate of the learning rate, $\hat{\omega}_n$.

Generate bootstrap samples $U_1^{(b)}, \dots, U_m^{(b)}$ and $V_1^{(b)}, \dots, V_n^{(b)}$, for $b = 1, \dots, B$.

Initialize $\omega^{(1)}$ and set $t = 1$.

repeat

$\omega = \omega^{(t)}$;

for b in $1 \dots B$ **do**

 Calculate $\text{HPD}_{\omega}^{(b)}$, the $100(1 - \alpha)\%$ highest Gibbs posterior density credible interval, with learning rate ω , based on the b^{th} bootstrap sample.

end

 Estimate the coverage probability $\hat{c}_{\alpha}(\omega) = B^{-1} |\{b : \text{HPD}_{\omega}^{(b)} \ni \hat{\theta}_{m,n}\}|$;

 Set $\Delta = \hat{c}_{\alpha}(\omega) - (1 - \alpha)$; Update $\omega^{(t+1)} = \omega + \kappa_t \Delta$; Set $t = t + 1$;

until $|\Delta| < \varepsilon$;

Return $\hat{\omega}_n = \omega^{(t)}$.

4 Numerical examples

4.1 Simulation studies

Since the AUC is invariant when random variables U and V undergo the same monotone increasing transformation, we fix the distribution of V to be standard normal and consider four examples for the distribution of U :

Example 1. $U \sim \mathbf{N}(2, 1)$ and $\theta^* = 0.9214$;

Example 2. $U \sim \text{SN}(3, 1, -4)$ —skew normal—and $\theta^* = 0.9665$;

Example 3. $U \sim 0.2\mathbf{N}(-1, 1) + 0.8\mathbf{N}(2, 0.5^2)$ and $\theta^* = 0.8185$;

Example 4. $U \sim 2 - \text{Exp}(1)$ and $\theta^* = 0.7895$.

Figure 2 provides a visualization of the two densities in each of the four examples. Note that these four examples capture binormality, a slight violation of binormality, a bimodal case, and one where U and V have different supports.

Here we compare four methods for inference on the AUC, including the Gibbs posterior, the rank-likelihood-based methods (BRL) described in Section 2.1, and another two Bayesian nonparametric methods. One is that presented in Gu et al. (2008) where the Bayesian bootstrap (BB) is used to approximate the posterior distribution of the ROC curve, and the marginal posterior for the AUC is obtained by calculating the area under each ROC curve sampled from its posterior. The other is based on the proposal in de Carvalho et al. (2013), where the distributions of U and V are each modeled as location-scale mixtures of Gaussians with a Dirichlet process prior on the mixing distribution. We use the default settings in the R package `dirichletprocess` (Ross and Markwick 2018) to get posterior samples from these Dirichlet process mixture (DPM) models. Using this,

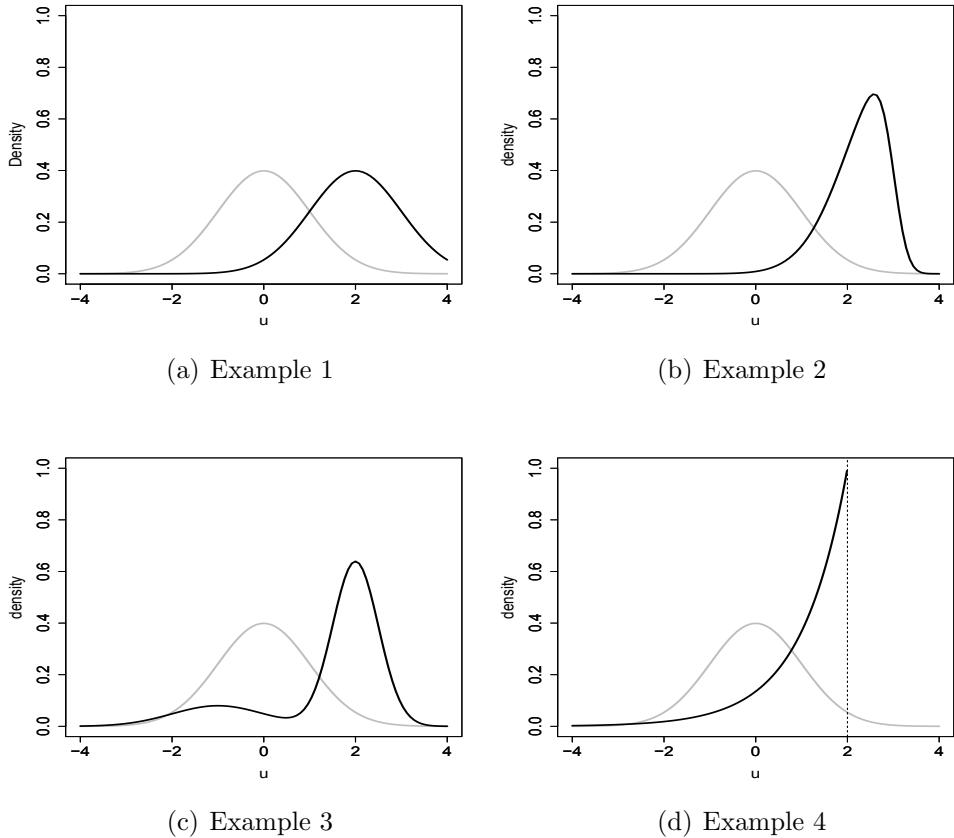


Figure 2: Density for U (black line) and the standard normal density for V (gray line) in Examples 1–4.

we can sample from the posterior predictive distributions of U and V , respectively, from which we readily obtain samples from the marginal AUC posterior.

For the Gibbs posterior, we use flat prior and follow Algorithm 1, where $B = 1000$ bootstrap samples are generated and $\kappa_t = (t + 1)^{-0.51}$ is set to satisfy (7). For the BRL and DPM, 50,000 posterior samples are drawn via MCMC, with burn-in of 10,000. For BB, we set the number of bootstrap samples to be 1000. With each increasing observation sizes $m = n \in \{25, 50, 75, 100, 125\}$, we take 1,000 data sets from the model as stated in the examples. Then for each data replication, we produce the AUC posterior for the four methods. We summarize the methods according to the following three metrics: the average absolute bias, computing by subtracting the true AUC from the posterior mean, the average posterior standard deviation, and the coverage probability of the 95% posterior credible intervals. The results are presented in Tables 1–4.

Both the Gibbs and BRL posteriors concentrate around the true AUC, but the former—thanks to its built-in robustness—tends to have a smaller bias than the latter. Furthermore, for large samples, the 95% credible intervals from the Gibbs posterior have coverage near the target level 0.95, while the corresponding BRL credible interval tend to under-cover, even in Example 1 where the binormality holds. Such a result is also demonstrated in Gu and Ghosal (2009). A possible explanation is that the posterior mean of BRL converges to θ^* but at a slower speed than the vanishing posterior spread.

n	abs(Bias)				Standard Deviation				Coverage Probability			
	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM
25	0.002	0.016	0.001	0.035	0.035	0.043	0.036	0.046	0.902	0.972	0.913	0.938
50	0.000	0.007	0.000	0.026	0.026	0.026	0.026	0.039	0.922	0.931	0.935	0.974
75	0.000	0.003	0.000	0.022	0.021	0.020	0.021	0.036	0.939	0.894	0.939	0.986
100	0.000	0.003	0.000	0.020	0.018	0.016	0.018	0.034	0.935	0.879	0.938	0.989
125	0.001	0.010	0.001	0.018	0.017	0.014	0.016	0.033	0.940	0.857	0.946	1.000

Table 1: Gibbs posterior compared with BRL, BB, and DPM for Example 1.

n	abs(Bias)				Standard Deviation				Coverage Probability			
	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM
25	0.005	0.022	0.001	0.026	0.020	0.035	0.021	0.033	0.997	0.949	0.857	0.958
50	0.001	0.006	0.000	0.018	0.015	0.017	0.015	0.027	0.912	0.904	0.885	0.985
75	0.000	0.001	0.000	0.015	0.013	0.012	0.013	0.025	0.919	0.902	0.911	0.995
100	0.000	0.002	0.000	0.013	0.012	0.010	0.011	0.024	0.931	0.907	0.923	0.999
125	0.000	0.004	0.000	0.013	0.011	0.009	0.010	0.023	0.944	0.861	0.927	1.000

Table 2: Gibbs posterior compared with BRL, BB, and DPM for Example 2.

For the two nonparametric methods, we expect that there would be generally less risk of model misspecification bias, but potentially some loss of efficiency due to treating the full U and V distributions as unknown. This is precisely what we found with DPM, in particular, the posterior standard deviations tend to be relatively large, which leads to over-coverage of its credible intervals. This extra variability may be due to the posterior for the mixture model including too many components. The other nonparametric method BB, however, has performance comparable to that of the Gibbs posterior, except in Example 2 where it tends to under-cover. Although BB’s performance is comparable to Gibbs in these settings, there are still advantages to the Gibbs formulation. In particular, since the Gibbs posterior focuses directly on the AUC it would be possible to incorporate available prior information, which would improve its performance. BB, on the other hand, is “non-informative” by definition and, therefore, cannot incorporate prior information. In these simulations, we used a flat prior for AUC and, hence, did not take advantage of this feature. But in real applications, using such prior information when available will give the Gibbs formulation an upper hand over BB.

Finally, we investigate the learning rate estimates under the Gibbs setting. Figure 3 shows, for each of the four simulation examples, the oracle learning rate compared to those obtained from the two suggested learning rate selection methods, namely, that based on Algorithm 1 and that in (9) based on asymptotic approximations. We find that the two have similar asymptotic behavior, and both closely follow the oracle learning rate, hence both will produce Gibbs posteriors whose credible sets achieve the nominal frequentist coverage probability asymptotically. However, the estimates produced by Algorithm 1 have less variability than the asymptotic approximation, which is why the former is our recommended method. Note, also, that the slope of the red line is less than -1 , which, in the log scale, is agrees with the tolerable decay rate—slower than $(m \vee n)^{-1}$ —suggested by the general theory in Section 3.2.

n	abs(Bias)				Standard Deviation				Coverage Probability			
	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM
25	0.002	0.020	0.000	0.026	0.065	0.064	0.065	0.058	0.919	0.922	0.929	0.896
50	0.002	0.016	0.002	0.019	0.046	0.044	0.047	0.049	0.933	0.900	0.935	0.941
75	0.000	0.011	0.000	0.017	0.037	0.035	0.038	0.046	0.921	0.887	0.938	0.977
100	0.000	0.008	0.000	0.017	0.032	0.030	0.033	0.044	0.936	0.897	0.950	0.988
125	0.001	0.003	0.001	0.019	0.029	0.027	0.030	0.043	0.934	0.890	0.942	0.992

Table 3: Gibbs posterior compared with BRL, BB, and DPM for Example 3.

n	abs(Bias)				Standard Deviation				Coverage Probability			
	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM	Gibbs	BRL	BB	DPM
25	0.000	0.025	0.001	0.033	0.066	0.063	0.065	0.064	0.925	0.902	0.936	0.930
50	0.000	0.024	0.000	0.026	0.045	0.043	0.046	0.054	0.937	0.844	0.941	0.973
75	0.001	0.020	0.001	0.025	0.037	0.033	0.038	0.051	0.930	0.788	0.940	0.983
100	0.000	0.003	0.000	0.023	0.032	0.028	0.033	0.049	0.942	0.861	0.952	0.990
125	0.000	0.020	0.000	0.022	0.029	0.026	0.029	0.047	0.938	0.803	0.950	0.999

Table 4: Gibbs posterior compared with BRL, BB, and DPM for Example 4

4.2 Real data analysis

Data consisting of two pancreatic cancer biomarkers' serum measurements was published by Wieand et al. (1989); see, also, the R package `logcondens`. We consider a diagnostic test that utilizes one of these biomarkers, namely, a cancer antigen (CA-125), to distinguish the case group (Group 1) from the control group (Group 0). This diagnostic test is a simple binary classifier, where the score is serum measurement of CA-125. As described in Section 1, individuals with scores higher than some threshold are diagnosed as positive outcomes. In this case-control study, the observations include $m = 90$ subjects from the case group and $n = 51$ subjects from the control group.

Our focus is on estimating the AUC of this binary classifier, and Table 5 presents results yielded by the Gibbs posterior, BRL, BB, and DPM. The Gibbs posterior results are based on a flat prior, and the learning rate chosen according to Algorithm 1. BRL and DPM are based on 50000 MCMC samples and 10,000 burn-in, and the initial values in BRL are $(a, b) = (2, 2)$ which follows Gu and Ghosal (2009). For BB, we use 10,000 bootstrap samples. We observe that DPM produces quite different results, compared to the other methods. The Gibbs posterior and BB yield similar results, and they have estimates slightly larger than that from BRL, with comparable standard deviations. The BRL credible interval is slightly shorter than the Gibbs and BB intervals but, in light of the simulation results presented above, especially in the case of relatively large samples like considered here, it is likely that the BRL intervals are “too short,” while the Gibbs and BB intervals are not.

5 Conclusion

In certain applications, the parameters of interest can be defined as minimizers of an appropriate risk function, separate from any statistical model. In such cases, one can avoid

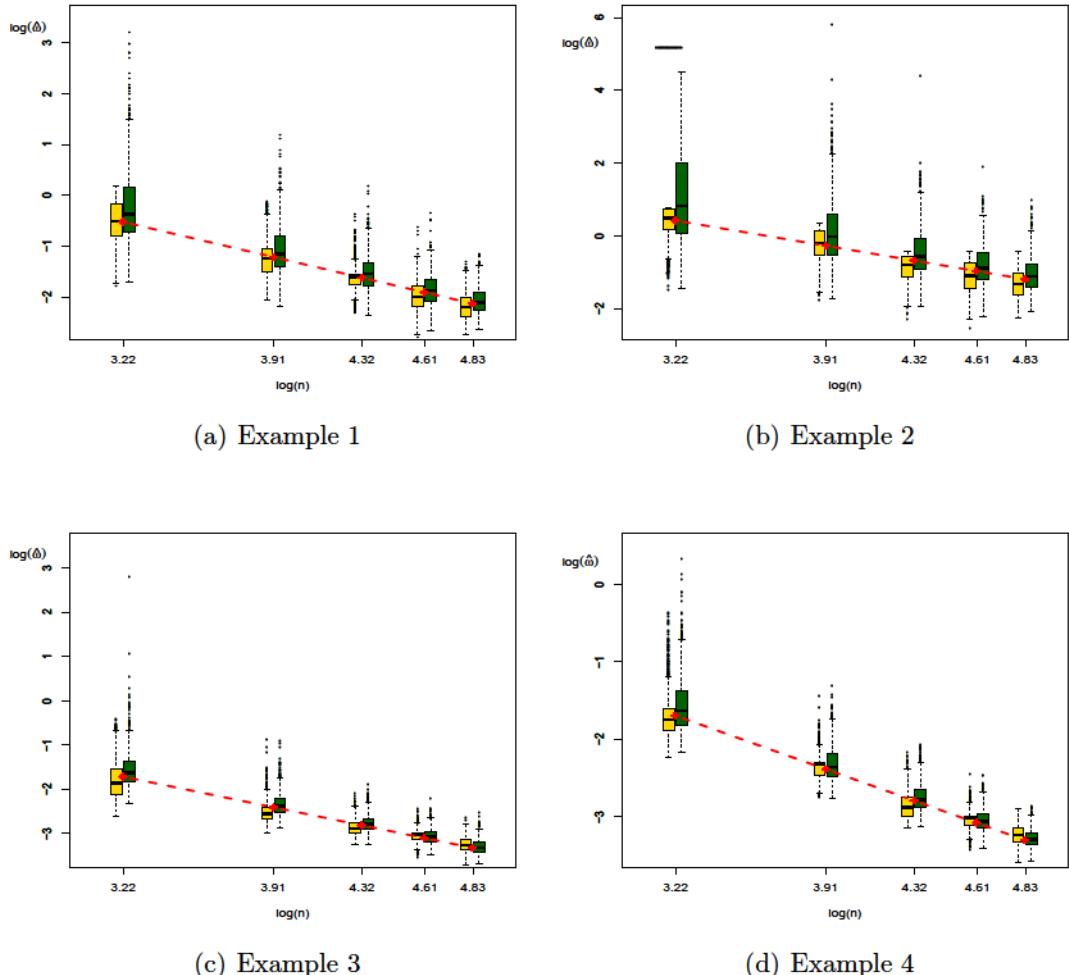


Figure 3: Boxplots of the learning rate estimates $\hat{\omega}$ from Algorithm 1 (yellow) and from (9) (green) versus sample size n , on the log scale, for Examples 1–4, with $m = n \in \{25, 50, 75, 100, 125\}$. The red dotted line represents logarithm of the oracle learning rate defined in Section 3.3.

potential model misspecification biases by working some kind of “model-free” approach. The present paper considered one such example, namely, inference on the AUC, where the state-of-the-art statistical model is one that depends on an infinite-dimensional nuisance parameter. As an alternative, we propose to construct a Gibbs posterior distribution for direct inference on the AUC, without specifying a model or introducing any nuisance parameters. This simplifies our computations and prior specifications, while allowing us to avoid potential model misspecification biases without sacrificing on the desirable asymptotic convergence properties. Moreover, a strategy for tuning the Gibbs posterior’s learning rate is recommended, that leads to credible intervals having the nominal frequentist coverage probability.

A direct extension of our work here is the inference on the analog of AUC in settings that involve three-group classifiers, namely, the volume under the ROC surface, or VUS (e.g., Mossman 1999). Similar to the set up here for the AUC, the VUS is defined as

	Gibbs	BRL	BB	DPM
Posterior mean	0.705	0.691	0.707	0.579
Posterior SD	0.045	0.046	0.047	0.054
Credible interval	(0.615, 0.795)	(0.598, 0.774)	(0.611, 0.793)	(0.470, 0.680)

Table 5: Posterior mean, standard deviation, and 95% credible intervals for the four methods—Gibbs, BRL, BB, and DPM—in the real-data example in Section 4.2.

$\mathbb{P}(T > U > V)$, where T is the score for the third group. Then much of the work presented here can be immediately generalized to the VUS case.

Acknowledgments

The authors thank the editors and anonymous reviewers for their helpful feedback on a previous version of the manuscript. This work is partially supported by the U.S. National Science Foundation, DMS-1811802.

A Proof of Theorem 1

Without loss of generality, we assume $n = m \wedge n$ and $n \rightarrow \infty$, which implies that $m = m_n \rightarrow \infty$ too. Next, when n (and, hence, m) is large, $\theta \mapsto \exp\{-\omega mn R_{m,n}(\theta)\}$ will blow up around $\theta = \hat{\theta}_{m,n}$ and, since the prior π is fixed—and positive in an interval containing θ^* and, hence, $\hat{\theta}_{m,n}$ —the Gibbs posterior will be dominated by the empirical risk term. Therefore, the prior does not affect the asymptotics so, for simplicity, we present the proof only for the case of a flat prior, $\pi(\theta) \equiv 1$.

By Chebyshev's inequality and the bias–variance decomposition of mean square error,

$$\Pi_{m,n}(\{\theta : |\theta - \theta^*| > K_n n^{-1/2}\}) \leq \frac{n}{K_n^2} \{V_{m,n} + (M_{m,n} - \theta^*)^2\}, \quad (10)$$

where $M_{m,n}$ and $V_{m,n}$ are the mean and variance of the Gibbs posterior distribution, respectively, and are given by

$$M_{m,n} = \hat{\theta}_{m,n} + \sigma_{m,n} \frac{\phi(A_{m,n}) - \phi(B_{m,n})}{\Phi(B_{m,n}) - \Phi(A_{m,n})}$$

$$V_{m,n} = \sigma_{m,n}^2 \left\{ 1 + \frac{A_{m,n}\phi(A_{m,n}) - B_{m,n}\phi(B_{m,n})}{\Phi(B_{m,n}) - \Phi(A_{m,n})} - \left[\frac{\phi(A_{m,n}) - \phi(B_{m,n})}{\Phi(B_{m,n}) - \Phi(A_{m,n})} \right]^2 \right\},$$

with ϕ and Φ the $\mathbb{N}(0, 1)$ density and distribution functions, respectively, and

$$A_{m,n} = -\mu_{m,n}\sigma_{m,n}^{-1} = -\hat{\theta}_{m,n}(2\omega mn)^{1/2}$$

$$B_{m,n} = (1 - \mu_{m,n})\sigma_{m,n}^{-1} = (1 - \hat{\theta}_{m,n})(2\omega mn)^{1/2}.$$

Since $\hat{\theta}_{m,n}$ is a consistent estimator of θ^* (see below), we clearly have that $A_{m,n} \rightarrow -\infty$ and $B_{m,n} \rightarrow \infty$, so those ratios involving ϕ and Φ above are all $O_p(1)$. Then we can

immediately conclude that $V_{m,n} = O_p((mn)^{-1})$ which takes care of the variance term. For the bias term, we first have that $\hat{\theta}_{m,n}$, the Mann–Whitney statistic, is an unbiased estimator of θ^* and its variance is upper-bounded by

$$\frac{\theta^*(1 - \theta^*)(m + n)}{mn}.$$

Therefore, for any $\varepsilon > 0$, there exists a number $L = L_\varepsilon$ such that

$$\mathsf{P}(n^{1/2}|\hat{\theta}_{m,n} - \theta^*| > L) \leq \varepsilon.$$

To see this, use Chebyshev's inequality and the bound on the variance of $\hat{\theta}_{m,n}$ to get that the left-hand side above is upper-bounded by

$$\frac{\theta^*(1 - \theta^*)(m + n)n}{L^2 mn}.$$

Since $(m + n)/m \leq 2$ and $\theta^*(1 - \theta^*) \leq 1/4$, we can take $L = L_\varepsilon$ sufficiently large that the previous display is less than ε . This implies that $|\hat{\theta}_{m,n} - \theta^*|$ and, hence, $|M_{m,n} - \theta^*|$ is $O_p(n^{-1/2})$. Putting everything together, we have that the right-hand side of (10) is

$$\frac{n}{K_n^2} \{O_p((mn)^{-1}) + O_p(n^{-1})\} = O_p(K_n^{-2}).$$

But since $K_n \rightarrow \infty$, we have that the upper-bound in (10) converges to 0 in P -probability as $(m, n) \rightarrow \infty$, proving the claim.

References

Alquier, P., Ridgway, J., and Chopin, N. (2016). On the properties of variational approximations of gibbs posteriors. *The Journal of Machine Learning Research*, 17(1):8374–8414.

Bamber, D. (1975). The area above the ordinal dominance graph and the area below the receiver operating characteristic graph. *Journal of Mathematical Psychology*, 12(4):387–415.

Bhattacharya, I. and Martin, R. (2020). Gibbs posterior inference on multivariate quantiles. Unpublished manuscript, [arXiv:2002.01052](https://arxiv.org/abs/2002.01052).

Bissiri, P. G., Holmes, C. C., and Walker, S. G. (2016). A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130.

Bradley, A. P. (1997). The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159.

Brodersen, K. H., Ong, C. S., Stephan, K. E., and Buhmann, J. M. (2010). The binormal assumption on precision-recall curves. In *2010 20th International Conference on Pattern Recognition*, pages 4263–4266. IEEE.

Cai, T. and Moskowitz, C. S. (2004). Semi-parametric estimation of the binormal ROC curve for a continuous diagnostic test. *Biostatistics*, 5(4):573–586.

Calì, C. and Longobardi, M. (2015). Some mathematical properties of the ROC curve and their applications. *Ricerche di Matematica*, 64(2):391–402.

Chernozhukov, V. and Hong, H. (2003). An MCMC approach to classical estimation. *J. Econometrics*, 115(2):293–346.

de Carvalho, V. I., Jara, A., Hanson, T. E., de Carvalho, M., et al. (2013). Bayesian nonparametric ROC regression modeling. *Bayesian Analysis*, 8(3):623–646.

Fasiolo, M., Goude, Y., Nedellec, R., and Wood, S. N. (2017). Fast calibrated additive quantile regression. *arXiv preprint arXiv:1707.03307*.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874.

Fraser, D. A. (2011). Is Bayes posterior just quick and dirty confidence? *Statistical Science*, 26(3):299–316.

Ghosal, S. and van der Vaart, A. (2017). *Fundamentals of Nonparametric Bayesian Inference*, volume 44 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.

Goddard, M. and Hinberg, I. (1990). Receiver operator characteristic (ROC) curves and non-normal data: an empirical study. *Statistics in Medicine*, 9(3):325–337.

Green, D. M. and Swets, J. A. (1966). *Signal detection theory and psychophysics*, volume 1. Wiley New York.

Grünwald, P. (2012). The safe Bayesian. In *International Conference on Algorithmic Learning Theory*, pages 169–183. Springer.

Grünwald, P. and Van Ommen, T. (2017). Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4):1069–1103.

Gu, J. and Ghosal, S. (2009). Bayesian ROC curve estimation under binormality using a rank likelihood. *Journal of Statistical Planning and Inference*, 139(6):2076–2083.

Gu, J., Ghosal, S., and Roy, A. (2008). Bayesian bootstrap estimation of ROC curve. *Statistics in Medicine*, 27(26):5407–5420.

Guignard, P. and Salehi, N. (1983). Validity of the Gaussian assumption in the analysis of ROC data obtained from scintigraphic-like images. *Physics in Medicine & Biology*, 28(12):1409.

Hanley, J. A. (1988). The robustness of the ‘binormal’ assumptions used in fitting ROC curves. *Medical Decision Making*, 8(3):197–203.

Hanley, J. A. (1989). Receiver operating characteristic (ROC) methodology: the state of the art. *Critical Reviews Diagnostic Imaging*, 29(3):307–335.

Hanley, J. A. and McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36.

Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 09(3):293–325.

Holmes, C. and Walker, S. (2017). Assigning a value to a power likelihood in a general Bayesian model. *Biometrika*, 104(2):497–503.

Huang, J. and Ling, C. X. (2005). Using AUC and accuracy in evaluating learning algorithms. *IEEE Transactions on knowledge and Data Engineering*, 17(3):299–310.

Jiang, W. and Tanner, M. A. (2008). Gibbs posterior for variable selection in high-dimensional classification and data mining. *The Annals of Statistics*, 36(5):2207–2231.

Kleijn, B. J. and van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *The Annals of Statistics*, 34(2):837–877.

Kleijn, B. J. and van der Vaart, A. W. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381.

Korolyuk, V. S. and Borovskich, Y. V. (2013). *Theory of U-statistics*, volume 273. Springer Science & Business Media.

Lusted, L. B. (1960). Logical analysis in roentgen diagnosis: memorial fund lecture. *Radiology*, 74(2):178–193.

Lyddon, S., Holmes, C., and Walker, S. (2019). General Bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106(2):465–478.

Macskassy, S. and Provost, F. (2004). Confidence bands for roc curves: Methods and an empirical study. Proceedings of the First Workshop on ROC Analysis in AI. August 2004.

Mann, H. B. and Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, pages 50–60.

Martin, R. (2019). False confidence, non-additive beliefs, and valid statistical inference. *arXiv preprint arXiv:1607.05051*.

Metz, C. E., Herman, B. A., and Shen, J.-H. (1998). Maximum likelihood estimation of receiver operating characteristic (ROC) curves from continuously-distributed data. *Statistics in Medicine*, 17(9):1033–1053.

Mossman, D. (1999). Three-way ROCs. *Medical Decision Making*, 19(1):78–89.

Provost, F. J. and Fawcett, T. (1997). Analysis and visualization of classifier performance: Comparison under imprecise class and cost distributions. In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining*, volume 97, pages 43–48.

Provost, F. J., Fawcett, T., and Kohavi, R. (1998). The case against accuracy estimation for comparing induction algorithms. In *Proceedings of the Fifteenth International Conference on Machine Learning*, volume 98, pages 445–453.

Ridgway, J., Alquier, P., Chopin, N., and Liang, F. (2014). PAC-Bayesian AUC classification and scoring. In *Advances in Neural Information Processing Systems*, pages 658–666.

Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407.

Ross, G. J. and Markwick, D. (2018). *dirichletprocess*: An R Package for Fitting Complex Bayesian Nonparametric Models.

Spackman, K. A. (1989). Signal detection theory: Valuable tools for evaluating inductive learning. In *Proceedings of the Sixth International Workshop on Machine Learning*, pages 160–163. Elsevier.

Swets, J. and Pickett, R. (1982). *Evaluation of Diagnostic Systems: Methods from Signal Detection Theory*. Academic Press series in cognition and perception. Academic Press.

Swets, J. A. (1973). The relative operating characteristic in psychology: a technique for isolating effects of response bias finds wide use in the study of perception and cognition. *Science*, 182(4116):990–1000.

Swets, J. A. (1986). Indices of discrimination or diagnostic accuracy: their ROCs and implied models. *Psychological Bulletin*, 99(1):100.

Syring, N. and Martin, R. (2017). Gibbs posterior inference on the minimum clinically important difference. *Journal of Statistical Planning and Inference*, 187:67–77.

Syring, N. and Martin, R. (2019a). Calibrating general posterior credible regions. *Biometrika*, 106:479–486.

Syring, N. and Martin, R. (2019b). Robust and rate optimal gibbs posterior inference on the boundary of a noisy image. *The Annals of Statistics*, to appear, [arXiv:1606.08400](https://arxiv.org/abs/1606.08400).

Syring, N. and Martin, R. (2020). Gibbs posterior concentration rates. In preparation.

Wieand, S., Gail, M. H., James, B. R., and James, K. L. (1989). A family of nonparametric statistics for comparing diagnostic markers with paired or unpaired data. *Biometrika*, 76(3):585–592.

Zhang, T. (2006a). From ϵ -entropy to KL-entropy: Analysis of minimum information complexity density estimation. *The Annals of Statistics*, 34(5):2180–2210.

Zhang, T. (2006b). Information-theoretic upper and lower bounds for statistical estimation. *IEEE Transactions on Information Theory*, 52(4):1307–1321.

Zou, K. H. and Hall, W. (2000). Two transformation models for estimating an ROC curve derived from continuous data. *Journal of Applied Statistics*, 27(5):621–631.