

ROBUST AND RATE-OPTIMAL GIBBS POSTERIOR INFERENCE ON THE BOUNDARY OF A NOISY IMAGE

BY NICHOLAS SYRING AND RYAN MARTIN

Washington University in St. Louis and North Carolina State University

Detection of an image boundary when the pixel intensities are measured with noise is an important problem in image segmentation. From a statistical point of view, a challenge is that likelihood-based methods require modeling the pixel intensities inside and outside the image boundary, even though these distributions are typically not of interest. Since misspecification of the pixel intensity distributions can negatively affect inference on the image boundary, it would be desirable to avoid this modeling step altogether. Towards this, we develop a robust Gibbsian approach that constructs a posterior distribution for the image boundary directly, without modeling the pixel intensities. We prove that the Gibbs posterior concentrates asymptotically at the minimax optimal rate, adaptive to the boundary smoothness. Monte Carlo computation of the Gibbs posterior is straightforward, and simulation results show that the corresponding inference is more accurate than that based on existing Bayesian methodology.

1. Introduction. In image analysis, the boundary or edge of the image is one of the most important features of the image, and extraction of this boundary is a critical step. An image consists of pixel locations and intensity values at each pixel, and the boundary can be thought of as a curve separating pixels of higher intensity from those of lower intensity. Applications of boundary detection are wide-ranging, e.g., [Martin, Fowlkes and Malik \(2004\)](#) use boundary detection to identify important features in pictures of natural settings, [Li et al. \(2010\)](#) identifies boundaries in medical images, and in [Yuan et al. \(2016\)](#) boundary detection helps classify the type and severity of wear on machines. For images with noiseless intensity, boundary detection has received considerable attention in the applied mathematics and computer science literature; see, e.g., [Ziou and Tabbone \(1998\)](#), [Maini and Aggarwal \(2009\)](#), [Li et al. \(2010\)](#), and [Anam, Uchino and Suetake \(2013\)](#). However, these approaches suffer from a number of difficulties. First, they can produce an estimate of the image boundary, but do not quantify estimation uncertainty. Second, these methods use a two-stage approach where

MSC 2010 subject classifications: Primary 62G20; secondary 62G05

Keywords and phrases: Adaptation, boundary detection, likelihood-free inference, model misspecification, posterior concentration rate

the image is first smoothed to filter out noise and then a boundary is estimated based on a calculated intensity gradient. This two-stage approach makes theoretical analysis of the method challenging, and no convergence results are presently known to the authors. Third, in our examples, these methods perform poorly on noisy data, and we suspect one reason for this is that the intensity gradient is less informative for the boundary when we observe noisy data. In the statistics literature, [Gu, Pati and Dunson \(2014\)](#) take a Bayesian approach to boundary detection and emphasize borrowing information to recover boundaries of similar objects in an image. Boundary detection using wombling is also a popular approach; see [Liang, Banerjee and Carlin \(2009\)](#), with applications to geography ([Lu and Carlin, 2004](#)), public health ([Ma and Carlin, 2007](#)), and ecology ([C. et al., 2010](#)). However, these techniques are used with areal or spatially aggregated data and are not suitable for the pixel data encountered in image analysis.

In Section 2, we give a detailed description of the image boundary problem, following the setup in [Li and Ghosal \(2017\)](#). They take a fully Bayesian approach, modeling the probability distributions of the pixel intensities both inside and outside the image. This approach is challenging because it often introduces nuisance parameters in addition to the image boundary. Therefore, in Section 3, we propose to use a so-called *Gibbs model*, where a suitable loss function is used to connect the data to the image boundary directly, rather than a probability model (e.g., [Catoni, 1997](#); [Zhang, 2006](#); [Jiang and Tanner, 2008](#); [Syring and Martin, 2017](#); [Bissiri, Holmes and Walker, 2016](#)).

We investigate the asymptotic convergence properties of the Gibbs posterior in Section 4. There we show that, if the boundary is α -Hölder smooth, then the Gibbs posterior concentrates around the true boundary at the rate $\{(\log n)/n\}^{\alpha/(\alpha+1)}$, which is minimax optimal ([Mammen and Tsybakov, 1995](#)) up to the logarithmic factor, relative to neighborhoods of the true boundary measured by the Lebesgue measure of a symmetric difference. Moreover, as a consequence of appropriately mixing over the number of knots in the prior, this rate is adaptive to the unknown smoothness α . To our knowledge, this is the first Gibbs posterior convergence rate result for an infinite-dimensional parameter, so the proof techniques used herein may be of general interest. Further, since the Gibbs posterior concentrates at the optimal rate without requiring a model for the pixel intensities, we claim that the inference on the image boundary is robust.

Computation of the Gibbs posterior is relatively straightforward and, in Section 5, we present a reversible jump Markov chain Monte Carlo method; R code to implement to the proposed Gibbs posterior inference is available at <https://github.com/nasyring/GibbsImage>. A comparison of inference

based on the proposed Gibbs posterior to that based on the fully Bayes approach in [Li and Ghosal \(2017\)](#) is shown in Section 6. For smooth boundaries, the two methods perform similarly, providing very accurate estimation. However, the Gibbs posterior is easier to compute, thanks to there being no nuisance parameters, and is notably more accurate than the Bayes approach when the model is misspecified or when the boundary is not everywhere smooth. Proofs of technical results are deferred to Appendix A.

2. Problem formulation. Let $\Omega \subset \mathbb{R}^2$ be a bounded region that represents the frame of the image; typically, Ω will be a square, say, $[-\frac{1}{2}, \frac{1}{2}]^2$, but, generally, we assume only that Ω is scaled to have unit Lebesgue measure. Data consists of pairs (X_i, Y_i) , $i = 1, \dots, n$, where X_i is a pixel location in Ω and Y_i is an intensity measurement at that pixel. The range of Y_i is context-dependent, and we consider both binary and real-valued cases below. The model asserts that there is a region $\Gamma \subset \Omega$ such that the intensity distribution is different depending on whether the pixel is inside or outside Γ . We consider the following model for the joint distribution P_Γ of pixel location and intensity, (X, Y) :

$$(1) \quad \begin{aligned} X &\sim g(x), \\ Y \mid (X = x) &\sim f_\Gamma(y) 1(x \in \Gamma) + f_{\Gamma^c}(y) 1(x \in \Gamma^c), \end{aligned}$$

where g is a density on Ω , f_Γ and f_{Γ^c} are densities on the intensity space, F_Γ and F_{Γ^c} are their respective distribution functions, and $1(\cdot)$ denotes an indicator function. That is, given the pixel location $X = x$, the distribution of the pixel intensity Y depends only on whether x is in Γ or Γ^c . Of course, more complex models are possible, e.g., where the pixel intensity distribution depends on the specific pixel location, but we will not consider such models here. We assume that there is a true, star-shaped set of pixels, denoted by Γ^* , with a known reference point in its interior. That is, any point in Γ^* can be connected to the reference point by a line segment fully contained in Γ^* . The observations $\{(X_i, Y_i) : i = 1, \dots, n\}$ are iid samples from P_{Γ^*} , and the goal is to make inference on Γ^* or, equivalently, its boundary $\gamma^* = \partial\Gamma^*$.

The density g for the pixel locations is of no interest and is taken to be known. The question is how to handle the two conditional distributions, f_Γ and f_{Γ^c} . [Li and Ghosal \(2017\)](#) take a fully Bayesian approach, modeling both f_Γ and f_{Γ^c} . By specifying these parametric models, they are obligated to introduce priors and carry out posterior computations for the corresponding parameters. Besides the efforts needed to specify models and priors and to carry out posterior computations, there is also a concern that the pixel intensity models might be misspecified, potentially biasing the inference on

Γ^* . Since the forms of f_Γ and f_{Γ^c} , as well as any associated parameters, are irrelevant to the boundary detection problem, it is natural to ask if inference can be carried out robustly, without modeling the pixel intensities.

We answer this question affirmatively, developing a Gibbs model for Γ in Section 3. In the present context, suppose we have a loss function $\ell_\Gamma(x, y)$ that measures how well an observed pixel location–intensity pair (x, y) agrees with a particular region Γ . The defining characteristic of ℓ_Γ is that Γ^* should be the unique minimizer of $\Gamma \mapsto R(\Gamma)$, where $R(\Gamma) = P_{\Gamma^*} \ell_\Gamma$ is the risk, i.e., the expected loss. A main contribution here, in Section 3.1, is specification of a loss function that meets this criterion. A necessary condition to construct such a loss function is that the distribution functions F_Γ and F_{Γ^c} are stochastically ordered. Imagine a gray-scale image; then, stochastic ordering means this image is lighter, on average, inside the boundary than outside the boundary, or vice-versa. In the specific context of binary images the pixel densities f_{Γ^*} and $f_{\Gamma^{*c}}$ are simply numbers between 0 and 1, and the stochastic ordering assumption means that, without loss of generality, $f_{\Gamma^*} > f_{\Gamma^{*c}}$, while for continuous images, again without loss of generality, $F_{\Gamma^*}(y) < F_{\Gamma^{*c}}(y)$ for all $y \in \mathbb{R}$. If we define the empirical risk,

$$(2) \quad R_n(\Gamma) = \frac{1}{n} \sum_{i=1}^n \ell_\Gamma(X_i, Y_i),$$

then, given a prior distribution Π for Γ , the Gibbs posterior distribution for Γ , denoted by Π_n , has the formula

$$(3) \quad \Pi_n(B) = \frac{\int_B e^{-nR_n(\Gamma)} \Pi(d\Gamma)}{\int e^{-nR_n(\Gamma)} \Pi(d\Gamma)},$$

where B is a generic Π -measurable set of Γ 's. Of course, (3) only makes sense if the denominator is finite; in Section 3, our risk functions R_n are non-negative so this integrability condition holds automatically.

Proper scaling of the loss in the Gibbs model is important (e.g., [Bissiri, Holmes and Walker, 2016](#); [Syring and Martin, 2018](#)), and here we will provide some context-specific scaling; see Sections 4 and 5.2. Our choice of prior Π for Γ is discussed in Section 3.2. Together, the loss function ℓ_Γ and the prior for Γ define the Gibbs model, no further modeling is required.

3. Gibbs model for the image boundary.

3.1. *Loss function.* To start, we consider inference on the image boundary when the pixel intensity is binary, i.e., $Y_i \in \{-1, +1\}$. In this case, the

densities, f_Γ and f_{Γ^c} , in (1) must be Bernoulli, so the likelihood is known. Even though the parametric form of the conditional distributions is known, the Gibbs approach only requires prior specification and posterior computation related to Γ , whereas the Bayes approach must also deal with the nuisance parameters in these Bernoulli conditional distributions. The binary case is relatively simple and will provide insights into how to formulate a Gibbs model in the more challenging continuous intensity problem.

A reasonable choice for the loss function ℓ_Γ is the following weighted misclassification error loss, depending on a parameter $h > 0$:

$$(4) \quad \ell_\Gamma(x, y) = h \mathbf{1}(y = +1, x \in \Gamma^c) + \mathbf{1}(y = -1, x \in \Gamma).$$

Note that putting $h = 1$ in (4) gives the usual misclassification error loss. In order for the Gibbs model to work the risk, or expected loss, must be minimized at the true Γ^* for some h . Picking h to ensure this property holds necessitates making a connection between the probability model in (1) and the loss in (4). The condition in (5) below is just the connection needed. With a slight abuse of notation, let f_{Γ^*} and $f_{\Gamma^{*c}}$ denote the conditional probabilities for the event $Y = +1$, given $X \in \Gamma^*$ and $X \in \Gamma^{*c}$, respectively. Recall our stochastic ordering assumption implies $f_{\Gamma^*} > f_{\Gamma^{*c}}$.

PROPOSITION 1. *Using the notation in the previous paragraph, if h is such that*

$$(5) \quad f_{\Gamma^*} \geq \frac{1}{1+h} \geq f_{\Gamma^{*c}},$$

then the risk function $R(\Gamma) = P_{\Gamma^} \ell_\Gamma$ is minimized at Γ^* .*

The condition in (5) deserves some further explanation. For example, if we know $f_{\Gamma^*} \geq \frac{1}{2} > f_{\Gamma^{*c}}$, then we take $h = 1$, which means that in (4) we penalize both intensities of 1 outside Γ and intensities of -1 inside Γ by a loss of 1. If, however, we know the overall image brightness is higher so that $f_{\Gamma^*} \geq \frac{4}{5} > f_{\Gamma^{*c}}$ then we take $h = 1/4$ in (4) and penalize bright pixels outside Γ by less than dull pixels inside Γ . To see why this loss balancing is so crucial, suppose the second case above holds so that $f_{\Gamma^*} = 4/5$ and $f_{\Gamma^{*c}} = 3/4$, but we take $h = 1$ anyway. Then, in (4), $\mathbf{1}(y = +1, x \in \Gamma^c)$ is very often equal to 1 while $\mathbf{1}(y = -1, x \in \Gamma)$ is very often 0. We will likely minimize the expected loss then by incorrectly taking Γ to be all of Ω so that the first term in the loss vanishes. Knowing a working h corresponds to having some prior information about f_{Γ^*} and $f_{\Gamma^{*c}}$, but we can also use the data to estimate a good value of h ; see Section 5.2.

Next, for the continuous case, we assume that the pixel intensities take values in \mathbb{R} . Our strategy is to modify the misclassification error (4) by working with a suitably discretized intensity measurement, reminiscent of threshold modeling. In particular, consider the following version of the misclassification error, depending on parameters (c, k, z) , with $c, k > 0$:

$$(6) \quad \ell_\Gamma(x, y) = k \mathbf{1}(y > z, x \in \Gamma^c) + c \mathbf{1}(y \leq z, x \in \Gamma).$$

Again, we claim that, for suitable (c, k, z) , the risk function is minimized at Γ^* . Let F_Γ and F_{Γ^c} denote the distribution functions corresponding to the densities f_Γ and f_{Γ^c} in (1), respectively. Recall our stochastic ordering assumption implies $F_{\Gamma^*}(z) < F_{\Gamma^{*c}}(z)$.

PROPOSITION 2. *If (c, k, z) in (6) satisfies*

$$(7) \quad F_{\Gamma^*}(z) \leq \frac{k}{k + c} \leq F_{\Gamma^{*c}}(z),$$

then the risk function $R(\Gamma) = P_{\Gamma^} \ell_\Gamma$ is minimized at Γ^* .*

The parameters k and c in (6) determine the scale of the loss as mentioned in Section 1, while z determines an intensity cutoff. According to the loss in (6), if a given pixel is located at $x \in \Gamma^c$, and with intensity y larger than cutoff z , it will incur a loss of $k > 0$. This implies that the true image Γ^* can be identified by working with a suitable version of the loss (6). A similar condition to (7), see Assumption 1 in Section 4, says what scaling is needed in order for the Gibbs posterior to concentrate at the optimal rate. Although the conditions on the scaling all involve the unknown distribution P_{Γ^*} , a good choice of (c, k, z) can be made based on the data alone, without prior information, and we discuss this strategy in Section 5.2.

3.2. Prior specification. We specify a prior distribution for the boundary of the region Γ by first expressing the pixel locations x in polar coordinates (θ, r) , an angle and radius, where $\theta \in [0, 2\pi]$ and $r > 0$, relative to the reference point described in Section 2. Li and Ghosal (2017) tested the influence of the reference point in simulations and found it to have little influence on their results. Using polar coordinates the boundary of Γ can be determined by the parametric curve $(\theta, \gamma(\theta))$. We proceed to model this curve γ .

Whether one is taking a Bayesian or Gibbsian approach, a natural strategy to model the image boundary is to express γ as a linear combination of suitable basis functions, i.e., $\gamma(\theta) = \hat{\gamma}_{D,\beta}(\theta) = \sum_{j=1}^D \beta_j B_{j,D}(\theta)$. Li and Ghosal (2017) use the eigenfunctions of the squared exponential periodic

kernel as their basis functions. Here we consider a model based on free knot b-splines, where the basis functions are defined recursively as

$$B_{i,1}(\theta) = 1(\theta \in [t_i, t_{i+1}])$$

$$B_{i,k}(\theta) = \frac{\theta - t_i}{t_{i+k-1} - t_i} B_{i,k-1}(\theta) + \frac{t_{i+k} - \theta}{t_{i+k} - t_{i+1}} B_{i+1,k-1}(\theta),$$

where $\theta \in [0, 2\pi]$, t_{-2}, t_{-1}, t_0 and $t_{D+1}, t_{D+2}, t_{D+3}$ are outer knots, t_1, \dots, t_D are inner knots, and $\beta = (\beta_1, \dots, \beta_D) \in (\mathbb{R}^+)^D$ is the vector of coefficients. Note that we restrict the coefficient vector β to be positive because the function values $\gamma(\theta)$ measure the radius of a curve from the origin. In the simulations in Section 6, the coefficients β_2, \dots, β_D are free parameters, while β_1 is calculated deterministically to force the boundary to be closed, i.e. $\gamma(0) = \gamma(2\pi)$, and we require $t_1 = 0$ and $t_D = 2\pi$; all other inner knots are free. Our model based on the b-spline representation performs as well as the eigenfunctions used in Li and Ghosal (2017) for smooth boundaries, and a bit better for boundaries with corners; see Section 6.

Therefore, the boundary curve γ is parametrized by an integer D and a D -vector β . We introduce a prior Π on (D, β) hierarchically as follows: D has a Poisson distribution with rate μ_D and, given D , the coordinates β_1, \dots, β_D of β are iid exponential with rate μ_β . These choices satisfy the technical conditions on Π detailed in Section 4. In our numerical experiments in Section 6, we take $\mu_D = 12$ and $\mu_\beta = 10$.

4. Gibbs posterior convergence. The Gibbs model depends on two inputs, namely, the prior and the loss function. In order to ensure that the Gibbs posterior enjoys desirable asymptotic properties, some conditions on both of these inputs are required. The first assumption listed below concerns the loss; the second concerns the true image boundary $\gamma^* = \partial\Gamma^*$; and the third concerns the prior. Here we will focus on the continuous intensity case, since the only difference between this and the binary case is that the latter provides the discretization for us.

ASSUMPTION 1. *Loss function parameters (c, k, z) in (6) satisfy*

$$(8) \quad F_{\Gamma^*}(z) < \frac{e^k - 1}{e^{c+k} - 1} \quad \text{and} \quad F_{\Gamma^{*c}}(z) > \frac{e^k - 1}{e^k - e^{-c}}.$$

Compared to the condition (7) that was enough to allow the loss function to identify the true Γ^* , condition (8) in Assumption 1 is only slightly stronger. This can be seen from the following inequality:

$$\frac{e^k - 1}{e^k - e^{-c}} > \frac{k}{k + c} > \frac{e^k - 1}{e^{c+k} - 1}.$$

However, if (c, k) are small, then the three quantities in the above display are all approximately equal, so Assumption 1 is not much stronger than what is needed to identify Γ^* . Again, these conditions on (c, k, z) can be understood as providing a meaningful scale to the loss function. Intuitively, the scale of the loss between observations receiving no loss versus some loss, expressed by parameters k and c , should be related to the level of information in the data. When $F_{\Gamma^*}(z)$ and $F_{\Gamma^{*c}}(z)$ are far apart, the data can more easily distinguish between F_{Γ^*} and $F_{\Gamma^{*c}}$, so we are free to assign larger losses than when $F_{\Gamma^*}(z)$ and $F_{\Gamma^{*c}}(z)$ are close and the data are relatively less informative.

The ability of a statistical method to make inference on the image boundary will depend on the true boundary smoothness. [Li and Ghosal \(2017\)](#) interpret γ^* as a function from the unit circle to the positive real line, and they formulate a Hölder smoothness condition for this function. Following the prior specification described in Section 3.2, we treat the boundary γ^* as a function from the interval $[0, 2\pi]$ to the positive reals, and formulate the smoothness condition on this arguably simpler version of the function. Since the reparametrization of the unit circle in terms of polar coordinates is smooth, it is easy to check that the Hölder smoothness condition (9) below is equivalent to that in [Li and Ghosal \(2017\)](#).

ASSUMPTION 2. *The true boundary function $\gamma^* : [0, 2\pi] \rightarrow \mathbb{R}^+$ is α -Hölder smooth, i.e., there exists a constant $L = L_{\gamma^*} > 0$ such that*

$$(9) \quad |(\gamma^*)^{([\alpha])}(\theta) - (\gamma^*)^{([\alpha])}(\theta')| \leq L|\theta - \theta'|^{\alpha - [\alpha]}, \quad \forall \theta, \theta' \in [0, 2\pi],$$

where $(\gamma^*)^{(k)}$ denotes the k^{th} derivative of γ^* and $[\alpha]$ denotes the largest integer less than or equal to α . Following the description of Γ^* in the introduction, we also assume that the reference point is strictly interior to Γ^* meaning that it is contained in an open set itself wholly contained in Γ^* so that γ^* is uniformly bounded away from zero. Moreover, the density g for X , as in (1), is uniformly bounded above by $\bar{g} := \sup_{x \in \Omega} g(x)$ and below by $\underline{g} := \inf_{x \in \Omega} g(x) \in (0, 1)$ on Ω .

General results are available on the error in approximating an α -Hölder smooth function by b-splines of the form specified in Section 3.2. Indeed, Theorem 6.10 in [Schumaker \(2007\)](#) implies that if γ^* satisfies (9), then

$$(10) \quad \forall d > 0, \exists \beta_d^* \in (\mathbb{R}^+)^d \text{ such that } \|\gamma^* - \hat{\gamma}_{d, \beta_d^*}\|_\infty \lesssim d^{-\alpha}.$$

Since $\gamma^*(\theta) > 0$, we can consider all coefficients to be positive, i.e., $\beta_d^* \in (\mathbb{R}^+)^d$; see Lemma 1(b) in [Shen and Ghosal \(2015\)](#). The next assumption about the prior makes use of the approximation property in (10).

ASSUMPTION 3. Let β_d^* , for $d > 0$, be as in (10). Then there exists $C, t > 0$ such that the prior Π for (D, β) satisfies, for all $d > 1$,

$$\begin{aligned} \log \Pi(D > d) &\lesssim -d \log d, \\ \log \Pi(D = d) &\gtrsim -d \log d, \\ \log \Pi(\|\beta - \beta_d^*\|_1 \leq kd^{-\alpha} \mid D = d) &\gtrsim -d \log\{1/(kd^{-\alpha})\}, \\ \log \Pi(\beta \notin [-m, m]^d \mid D = d) &\lesssim \log d - Cm^t. \end{aligned}$$

The first two conditions in Assumption 3 ensure that the prior on D is sufficiently spread out while the second two conditions ensure that there is sufficient prior support near β 's that approximate γ^* well. Our suggested Poisson prior for D and independent exponential prior for β , given D , satisfy these conditions; see also [Shen and Ghosal \(2015\)](#). Assumption 3 is also needed in [Li and Ghosal \(2017\)](#) for convergence of the Bayesian posterior at the optimal rate. However, the Bayes model also requires assumptions about the priors on the nuisance parameters, e.g., Assumption C in [Li and Ghosal \(2017\)](#), which are not necessary in our approach here.

In what follows, let $A \Delta B$ denote the symmetric difference of sets in Ω and $\lambda(A \Delta B)$ its Lebesgue measure.

THEOREM 1. *With a slight abuse of notation, let Π denote the prior for Γ , induced by that on (D, β) , and Π_n the corresponding Gibbs posterior (3). Under Assumptions 1–3, there exists a constant $M > 0$ such that*

$$P_{\Gamma^*} \Pi_n(\{\Gamma : \lambda(\Gamma^* \Delta \Gamma) > M\varepsilon_n\}) \rightarrow 0 \quad \text{as } n \rightarrow \infty,$$

where $\varepsilon_n = \varepsilon_n(\alpha) = \{(\log n)/n\}^{\alpha/(\alpha+1)}$ and α is the Hölder smoothness coefficient in Assumption 2.

Theorem 1 says that, as the sample size increases, the Gibbs posterior places its mass on a shrinking neighborhood of the true boundary γ^* . The rate, given by the size of the neighborhood, is optimal according to [Mammen and Tsybakov \(1995\)](#), up to a logarithmic factor, and adaptive since the prior does not depend on the unknown smoothness α .

5. Computation.

5.1. *Sampling algorithm.* We use reversible jump Markov chain Monte Carlo (e.g., [Green, 1995](#)) to sample from the Gibbs posterior. These methods have been used successfully in Bayesian free-knot spline regression problems; see, e.g., [Denison, Mallick and Smith \(1998\)](#) and [DiMatteo, Genovese and](#)

Kass (2001). Although sampling is more complicated when the number and locations of knots are random instead of fixed, the resulting spline functions can do a better job fitting curves with low smoothness.

To summarize the algorithm, we start with the prior distribution Π for (D, β) as discussed in Section 3.2. Next, we need to initialize values of D , the knot locations $\{t_{-2}, \dots, t_{D+3}\}$, and the values of β_2, \dots, β_D . The value of β_1 is then calculated numerically to force closure. We choose $D = 12$ with $t_{-2} = -2$, $t_{-1} = -1$, $t_0 = -0.5$, $t_{13} = 2\pi + 0.5$, $t_{14} = 2\pi + 1$, $t_{15} = 2\pi + 2$ and t_1, \dots, t_{12} evenly spaced in $[0, 2\pi]$. We set inner knots $t_0 = 0$ and $t_D = 2\pi$ while the other inner knot locations remain free to change in birth, death, and relocation moves; we also set $\beta_2 = \beta_3 = \dots = \beta_{12} = 0.1$. Then the following three steps constitutes a single iteration of the reversible jump Markov chain Monte Carlo algorithm to be repeated until the desired number of samples are obtained:

1. Use Metropolis-within-Gibbs steps to update the elements of the β vector, again solving for β_1 to force closure at the end. In our examples we use normal proposals centered at the current value of the element of the β vector, and with standard deviation 0.10.
2. Randomly choose to attempt either a birth, death, or relocation move to add a new inner knot, delete an existing inner knot, or move an inner knot.
3. Attempt the jump move proposed in Step 2. The β vector must be appropriately modified when adding or deleting a knot, and again we must solve for β_1 . Details on the calculation of acceptance probabilities for each move can be found in Denison, Mallick and Smith (1998) and DiMatteo, Genovese and Kass (2001).

R code to implement this Gibbs posterior sampling scheme, along with the empirical loss scaling method described in Section 5.2, is available at <https://github.com/nasyring/GibbsImage>.

5.2. *Loss scaling.* It is not clear how to select (c, k, z) to satisfy Assumption 1 without knowledge of F_{Γ^*} and $F_{\Gamma^{*c}}$. However, it is fairly straightforward to select values of (c, k, z) based on the data which are likely to meet the required condition. First, we need a notion of optimal (c, k, z) values. If we knew F_{Γ^*} and $F_{\Gamma^{*c}}$, then we would select z to maximize $F_{\Gamma^{*c}}(z) - F_{\Gamma^*}(z)$ because this choice of z gives us the point at which F_{Γ^*} and $F_{\Gamma^{*c}}$ are most easily distinguished. Then, we would choose (c, k) to be the largest values such that (8) holds. Intuitively, we want large values of (c, k) so that the loss function in (6) is more sensitive to departures from γ^* .

Since we do not know F_{Γ^*} and $F_{\Gamma^{*c}}$, we estimate $F_{\Gamma^*}(z)$ and $F_{\Gamma^{*c}}(z)$ from

the data. In order to do this, we need a rough estimate of γ^* to define the regions Γ^* and Γ^{*c} . Our approach is to model γ with a b-spline, as before, and estimate γ^* several times by minimizing (6) using several different values of the loss scaling parameters (c, k, z) . Specifically, set a grid of z values z_1, z_2, \dots, z_g , and for each z_j , find $(c, k) = (c_j, k_j)$ that satisfy

$$\frac{k_j}{k_j + c_j} = \frac{|\{i : y_i \leq z_j\}|}{n}.$$

Next, estimate γ^* by minimizing (6) using (c_j, k_j, z_j) . The estimate of γ^* provides estimates of the regions Γ^* and Γ^{*c} which we use to calculate the sample proportions

$$\hat{F}_{\Gamma^*} := \frac{|\{i : y_i \leq z_j, x_i \in \hat{\Gamma}^*\}|}{|\{i : x_i \in \hat{\Gamma}^*\}|} \quad \text{and} \quad \hat{F}_{\Gamma^{*c}} := \frac{|\{i : y_i \leq z_j, x_i \in \hat{\Gamma}^{*c}\}|}{|\{i : x_i \in \hat{\Gamma}^{*c}\}|}.$$

Then, the approximately optimal value is $z = \arg \max_{z_j} \hat{F}_{\Gamma^{*c}}(z_j) - \hat{F}_{\Gamma^*}(z_j)$. Finally, choose the approximately optimal values of (c, k) to satisfy (8) replacing $F_{\Gamma^*}(z)$ and $F_{\Gamma^{*c}}(z)$ by their estimates $\hat{F}_{\Gamma^*}(z)$ and $\hat{F}_{\Gamma^{*c}}(z)$.

Based on the simulations in Section 6, this method produces values of (c, k, z) very close to their optimal values. Importantly, the estimated (c, k) are more likely to be smaller than their optimal values than larger, which makes our estimates more likely to satisfy (8). This is a consequence of the stochastic ordering of F_{Γ^*} and $F_{\Gamma^{*c}}$. Unless the classifier we obtain by minimizing (6) is perfectly accurate, we will tend to mix together samples from F_{Γ^*} and $F_{\Gamma^{*c}}$ in our estimates. If we estimate $F_{\Gamma^*}(z)$ with some observations from F_{Γ^*} and some from $F_{\Gamma^{*c}}$, we will tend to overestimate $F_{\Gamma^*}(z)$, and vice versa we will tend to underestimate $F_{\Gamma^{*c}}(z)$. These errors will cause (c, k) to be underestimated, and therefore more likely to satisfy (8).

6. Numerical examples. We tested our Gibbs model on data from both binary and continuous images following much the same setup as in Li and Ghosal (2017). The pixel locations in $\Omega = [-\frac{1}{2}, \frac{1}{2}]^2$ are sampled by starting with a fixed $m \times m$ grid in Ω and making a small random uniform perturbation at each grid point. Several different pixel intensity distributions are considered. We consider two types of shapes for Γ^* : an ellipse with center $(0.1, 0.1)$, rotated at an angle of 60 degrees, with major axis length 0.35 and minor axis length 0.25; and a centered equilateral triangle of height 0.5. The ellipse boundary will test the sensitivity of the model to boundaries which are off-center while the triangle tests the model's ability to identify non-smooth boundaries.

We consider four binary and four continuous intensity images and compare with the Bayesian method of [Li and Ghosal \(2017\)](#) as implemented in the *BayesBD* package ([Syring and Li, 2017a](#)) available on CRAN and described in [Syring and Li \(2017b\)](#).

- B1. Ellipse image, $m = 100$, and F_{Γ^*} and $F_{\Gamma^{*c}}$ are Bernoulli with parameters 0.5 and 0.2, respectively.
- B2. Same as B1 but with triangle image.
- B3. Ellipse image, $m = 500$, and F_{Γ^*} and $F_{\Gamma^{*c}}$ are Bernoulli with parameters 0.25 and 0.2, respectively.
- B4. Same as B3 but with triangle image.
- C1. Ellipse image, $m = 100$, and F_{Γ^*} and $F_{\Gamma^{*c}}$ are $N(4, 1.5^2)$ and $N(1, 1)$, respectively.
- C2. Same as C1 but with triangle image.
- C3. Ellipse image, $m = 100$, and F_{Γ^*} and $F_{\Gamma^{*c}}$ are $0.2 N(2, 10) + 0.8 N(0, 1)$, a normal mixture, and $N(0, 5)$, respectively.
- C4. Ellipse image, $m = 100$, and F_{Γ^*} and $F_{\Gamma^{*c}}$ are t distributions with 3 degrees of freedom and non-centrality parameters 1 and 0, respectively.

For binary images, the likelihood must be Bernoulli, so the Bayesian model is correctly specified in cases B1–B4. For the continuous examples in C1–C4, we assume a Gaussian likelihood for the Bayesian model. Then, cases C1 and C2 will show whether or not the Gibbs model can compete with the Bayesian model when the model is correctly specified, while cases C3 and C4 will demonstrate the superiority of the Gibbs model over the Bayesian model under misspecification. Again, the Gibbs model has the added advantage of not having to specify priors for or sample values of the mean and variance associated with the normal conditional distributions.

We replicated each scenario 100 times for both the Gibbs and Bayesian models, each time producing a posterior sample of size 4000 after a burn in of 1000 samples. We recorded the errors—Lebesgue measure of the symmetric difference—for each run along with the estimated loss function parameters for the Gibbs models for continuous images. The results are summarized in Tables 1–2. We see that the Gibbs model is competitive with the fully Bayesian model in Examples B1–B4, C1, and C2, when the likelihood is correctly specified. When the likelihood is misspecified, the Bayesian model may fail, as in Examples C3 and C4. The Gibbs model does not depend on a likelihood, only the stochastic ordering, so it continues to perform well in these non-Gaussian examples. From Table 2, we see that the empirical method described in Section 5.2 is able to select parameters for the loss function in (6) close to the optimal values and meeting Assumption 1.

TABLE 1
Average errors (and standard deviations) for each example.

Model	B1	B2	B3	B4	C1	C2	C3	C4
Bayes	0.00 (0.00)	0.02 (0.00)	0.01 (0.00)	0.02 (0.01)	0.03 (0.03)	0.04 (0.03)	0.11 (0.06)	0.10 (0.05)
	0.01 (0.00)	0.01 (0.00)	0.02 (0.01)	0.02 (0.01)	0.01 (0.00)	0.01 (0.00)	0.01 (0.01)	0.01 (0.01)

TABLE 2
Average (and optimal) values of the parameters (c, k, z) in (6).

Parameter	C1	C2	C3	C4
c	1.45 (1.86)	1.47 (1.86)	0.80 (1.27)	0.71 (0.80)
	2.30 (2.36)	2.29 (2.36)	0.26 (0.34)	0.71 (0.75)
k	2.43 (2.40)	2.39 (2.40)	-1.83 (-1.76)	0.46 (0.46)

Figures 1 and 2 show the results of the Bayesian and Gibbs models for one simulation run in each of Examples B1–B2 and C1–C4, respectively. The 95% credible regions, as in [Li and Ghosal \(2017\)](#), are highlighted in gray around the posterior means. That is, let $u_i = \sup_{\theta} \{s(\theta)^{-1} |\gamma_i(\theta) - \hat{\gamma}(\theta)|\}$, where $\gamma_i(\theta)$ is the i^{th} posterior boundary sample, $\hat{\gamma}(\theta)$ is the pointwise posterior mean and $s(\theta)$ the pointwise standard deviation of the $\gamma(\theta)$ samples. If τ is the 95th percentile of the u_i 's, then a 95% uniform credible band is given by $\hat{\gamma}(\theta) \pm \tau s(\theta)$. The results of cases B2 and C2 suggest that free-knot b-splines may do a better job of approximating non-smooth boundaries than the kernel basis functions used by [Li and Ghosal \(2017\)](#). In particular, the reversible jump sampling method with its relocation moves allowed knots to move towards the corners of the triangle, thereby improving estimation of the boundary over b-splines with fixed knots.

Acknowledgments. The authors thank the anonymous reviewers for their helpful comments. This work is partially supported by the U. S. National Science Foundation, grant DMS–1811802.

APPENDIX A: PROOFS

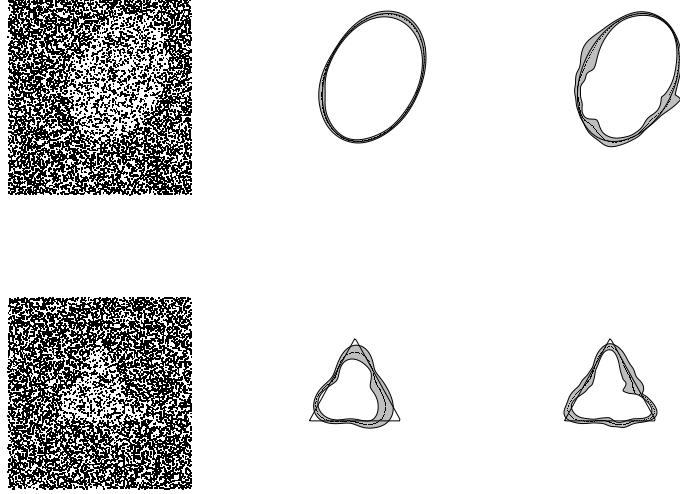


FIG 1. From top, Examples B1–B2. In each row, the observed image is on the left, the Bayesian posterior mean estimator (Li and Ghosal, 2017) is in the middle, and the Gibbs posterior mean estimator is on the right. Solid lines show the true image boundary, dashed lines are the estimates, and gray regions are 95% credible bands.

A.1. Proof of Proposition 1. By the definition of the loss function in (4), for a fixed h and for any $\Gamma \subset \Omega$, we have

$$\begin{aligned} \ell_\Gamma(x, y) - \ell_{\Gamma^*}(x, y) &= h 1(y = +1, x \in \Gamma^c) - h 1(y = +1, x \in \Gamma^{*c}) \\ &\quad + 1(y = -1, x \in \Gamma) - 1(y = -1, x \in \Gamma^*) \\ &= h 1(y = +1, x \in \Gamma^* \setminus \Gamma) - 1(y = -1, x \in \Gamma^* \setminus \Gamma) \\ &\quad + 1(y = -1, x \in \Gamma \setminus \Gamma^*) - h 1(y = +1, x \in \Gamma \setminus \Gamma^*). \end{aligned}$$

Then the expectation of the loss difference above is

$$P(X \in \Gamma^* \setminus \Gamma) (h f_{\Gamma^*} + f_{\Gamma^*} - 1) + P(X \in \Gamma \setminus \Gamma^*) (1 - f_{\Gamma^{*c}} - h f_{\Gamma^{*c}}),$$

where the probability statement is with respect to the density g of X . By Assumption 2, the density g is bounded away from zero on Ω , so the expectation of the loss difference is zero if and only if $\Gamma = \Gamma^*$. The expectation can also be lower bounded by

$$P(X \in \Gamma \Delta \Gamma^*) \min\{h f_{\Gamma^*} + f_{\Gamma^*} - 1, 1 - f_{\Gamma^{*c}} - h f_{\Gamma^{*c}}\}.$$

According to (5) in Proposition 1, both terms in the minimum are positive. Therefore, $R(\Gamma) \geq R(\Gamma^*)$ with equality if and only if $\Gamma = \Gamma^*$.

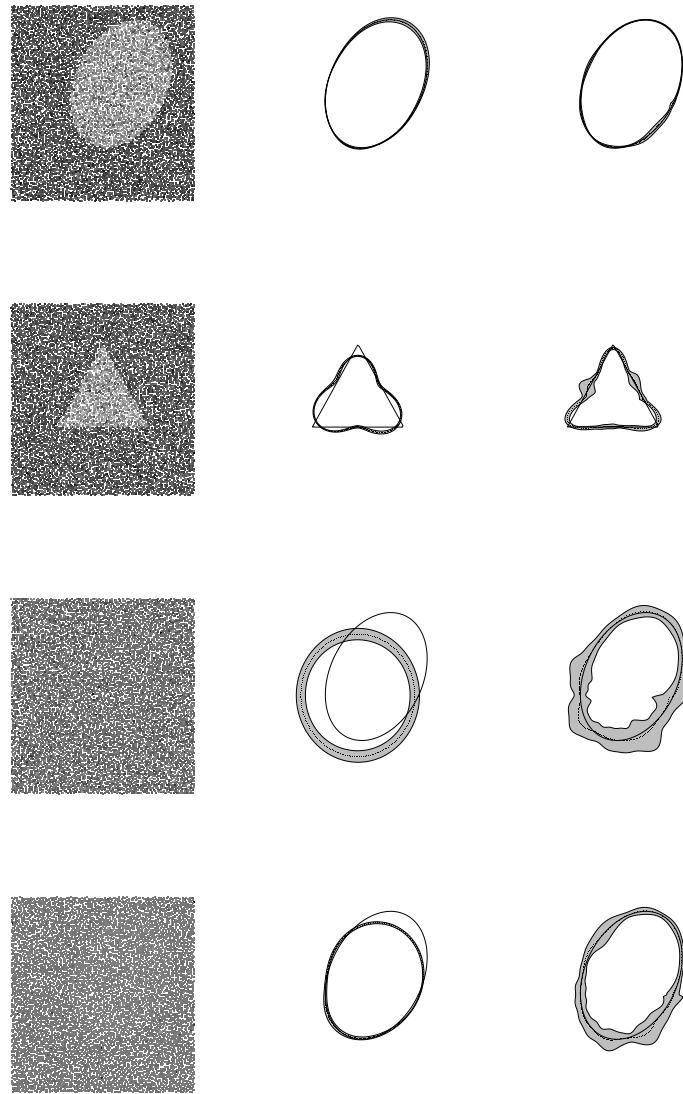


FIG 2. Same as Figure 1, but for Examples C1–C4.

A.2. Proof of Proposition 2. The proof here is very similar to that of Proposition 1. By the definition of the loss function in (6), for any fixed (c, k, z) and for any $\Gamma \subset \Omega$, we get

$$\begin{aligned}\ell_\Gamma(x, y) - \ell_{\Gamma^*}(x, y) &= k 1(y \geq z, x \in \Gamma^c) - k 1(y \geq z, x \in \Gamma^{*c}) \\ &\quad + c 1(y < z, x \in \Gamma) - c 1(y < z, x \in \Gamma^*) \\ &= k 1(y \geq z, x \in \Gamma^* \setminus \Gamma) - c 1(y < z, x \in \Gamma^* \setminus \Gamma) \\ &\quad + c 1(y < z, x \in \Gamma \setminus \Gamma^*) - k 1(y \geq z, x \in \Gamma \setminus \Gamma^*).\end{aligned}$$

Then, the expectation of the loss difference above is given by

$$P(X \in \Gamma^* \setminus \Gamma) \{k - kF_{\Gamma^*}(z) - cF_{\Gamma^*}(z)\} + P(X \in \Gamma \setminus \Gamma^*) \{cF_{\Gamma^{*c}}(z) - k + kF_{\Gamma^{*c}}(z)\},$$

where the probability statement is with respect to the density g of X . As in the proof of Proposition 1, this quantity is zero if and only if $\Gamma = \Gamma^*$. It can also be lower bounded by

$$P(X \in \Gamma \Delta \Gamma^*) \min\{k - kF_{\Gamma^*}(z) - cF_{\Gamma^*}(z), cF_{\Gamma^{*c}}(z) - k + kF_{\Gamma^{*c}}(z)\}.$$

Given the condition (7) in Proposition 2, both terms in the minimum are positive. Therefore, $R(\Gamma) \geq R(\Gamma^*)$ with equality if and only if $\Gamma = \Gamma^*$.

A.3. Preliminary results. Towards proving Theorem 1, we need several lemmas. The first draws a connection between the distance between defined by the Lebesgue measure of the symmetric difference and the sup-norm between the boundary functions.

LEMMA 1. *Suppose Γ^* , with boundary $\gamma^* = \partial\Gamma^*$, satisfies Assumption 2, in particular, $\underline{\gamma}^* := \inf_{\theta \in [0, 2\pi]} \gamma^*(\theta) > 0$. Take any $\Gamma \subset \Omega$, with $\gamma = \partial\Gamma$, such that $\lambda(\Gamma \Delta \Gamma^*) > \delta$ for some fixed $\delta > 0$, and any $\tilde{\Gamma} \subset \Omega$ such that $\tilde{\gamma} = \partial\tilde{\Gamma}$ satisfies $\|\tilde{\gamma} - \gamma\|_\infty < \omega\delta$, where $\omega \in (0, 1)$. Then*

$$\lambda(\tilde{\Gamma} \Delta \Gamma^*) > \frac{4\delta}{\underline{\gamma}^*} \left(\frac{1}{\text{diam}(\Omega)} - \pi\omega \right),$$

where $\text{diam}(\Omega) = \sup_{x, x' \in \Omega} \|x - x'\|$ is the diameter of Ω . So, if $\omega < \{\pi \text{diam}(\Omega)\}^{-1}$, then the lower bound is a positive multiple of δ .

PROOF. The first step is to connect the symmetric difference-based distance to the L_1 distance between boundary functions. A simple conversion

to polar coordinates gives

$$\begin{aligned}
\lambda(\Gamma \Delta \Gamma^*) &= \int_{\Gamma \Delta \Gamma^*} d\lambda \\
&= \int_0^{2\pi} \int_{\gamma(\theta) \wedge \gamma^*(\theta)}^{\gamma(\theta) \vee \gamma^*(\theta)} r dr d\theta \\
&= \frac{1}{2} \int_0^{2\pi} \{\gamma(\theta) \wedge \gamma(\theta^*)\}^2 - \{\gamma(\theta) \vee \gamma(\theta^*)\}^2 d\theta \\
&= \frac{1}{2} \int_0^{2\pi} |\gamma(\theta) - \gamma^*(\theta)| |\gamma(\theta) + \gamma^*(\theta)| d\theta.
\end{aligned}$$

If we let $\underline{\gamma}^* = \inf_\theta \gamma^*(\theta)$, then it is easy to verify that

$$\underline{\gamma}^* \leq |\gamma(\theta) + \gamma^*(\theta)| \leq \text{diam}(\Omega), \quad \forall \theta \in [0, 2\pi].$$

Therefore,

$$(11) \quad \frac{1}{2} \underline{\gamma}^* \|\gamma - \gamma^*\|_1 \leq \lambda(\Gamma \Delta \Gamma^*) \leq \frac{1}{2} \text{diam}(\Omega) \|\gamma - \gamma^*\|_1.$$

Next, if $\lambda(\Gamma \Delta \Gamma^*) > \delta$, which is positive by Assumption 2, then it follows from the right-most inequality in (11) that $\text{diam}(\Omega) \|\gamma - \gamma^*\|_1 > 2\delta$ and, by the triangle inequality,

$$\text{diam}(\Omega) \{\|\gamma - \tilde{\gamma}\|_1 + \|\tilde{\gamma} - \gamma^*\|_1\} > 2\delta.$$

We have $\|\gamma - \tilde{\gamma}\|_1 \leq 2\pi \|\gamma - \tilde{\gamma}\|_\infty$ which, by assumption, is less than $2\pi\omega\delta$. Consequently,

$$\text{diam}(\Omega) \{2\pi\omega\delta + \|\tilde{\gamma} - \gamma^*\|_1\} > 2\delta$$

and, hence,

$$\|\tilde{\gamma} - \gamma^*\|_1 > \frac{2\delta}{\text{diam}(\Omega)} - 2\pi\omega\delta.$$

By the left-most inequality in (11), we get

$$\lambda(\tilde{\Gamma} \Delta \Gamma^*) > \frac{4\delta}{\underline{\gamma}^* \text{diam}(\Omega)} - \frac{4\pi\omega\delta}{\underline{\gamma}^*} = \frac{4\delta}{\underline{\gamma}^*} \left(\frac{1}{\text{diam}(\Omega)} - \pi\omega \right),$$

which is the desired bound. It follows immediately that the lower bound is a positive multiple of δ if $\omega < \{\pi \text{diam}(\Omega)\}^{-1}$. \square

The next lemma shows that we can control the expectation of the integrand in the Gibbs posterior under the condition (8) on the tuning parameters (c, k, z) in the loss function definition.

LEMMA 2. *If (8) holds, then $P_{\Gamma^*}e^{-(\ell_{\Gamma} - \ell_{\Gamma^*})} < 1 - \rho\lambda(\Gamma^* \Delta \Gamma)$ for a constant $\rho \in (0, 1)$.*

PROOF. From the proof of Proposition 2, we have

$$\begin{aligned} \ell_{\Gamma}(x, y) - \ell_{\Gamma^*}(x, y) &= k 1(y \geq z, x \in \Gamma^* \setminus \Gamma) - c 1(y < z, x \in \Gamma^* \setminus \Gamma) \\ &\quad + c 1(y < z, x \in \Gamma \setminus \Gamma^*) - k 1(y \geq z, x \in \Gamma \setminus \Gamma^*). \end{aligned}$$

The key observation is that, if $x \notin \Gamma \Delta \Gamma^*$, then the loss difference is 0 and, therefore, the exponential of the loss difference is 1. Taking expectation with respect P_{Γ^*} , we get

$$\begin{aligned} P_{\Gamma^*}e^{-(\ell_{\Gamma} - \ell_{\Gamma^*})} &= P_g(X \notin \Gamma \Delta \Gamma) \\ &\quad + \{e^{-k}(1 - F_{\Gamma^*}(z)) + e^c F_{\Gamma^*}(z)\} P_g(X \in \Gamma^* \setminus \Gamma) \\ &\quad + \{e^{-c} F_{\Gamma^{*c}}(z) + e^k - e^k F_{\Gamma^{*c}}(z)\} P_g(X \in \Gamma \setminus \Gamma^*). \end{aligned}$$

If we define

$$\kappa = \max\{e^{-k}(1 - F_{\Gamma^*}(z)) + e^c F_{\Gamma^*}(z), e^{-c} F_{\Gamma^{*c}}(z) + e^k - e^k F_{\Gamma^{*c}}(z)\},$$

then it follows from (8) that $\kappa < 1$ and

$$\begin{aligned} P_{\Gamma^*}e^{-(\ell_{\Gamma} - \ell_{\Gamma^*})} &\leq 1 - P_g(X \in \Gamma \Delta \Gamma^*) + \kappa P_g(X \in \Gamma \Delta \Gamma^*) \\ &= 1 - (1 - \kappa) P_g(X \in \Gamma \Delta \Gamma^*). \end{aligned}$$

Since $P_g(X \in \Gamma \Delta \Gamma^*) \geq g\lambda(\Gamma \Delta \Gamma^*)$, the claim holds with $\rho = (1 - \kappa)g$. \square

The next lemma yields a necessary lower bound on the denominator of the Gibbs posterior distribution. The neighborhoods G_n are simpler than those in Lemma 1 of [Shen and Wasserman \(2001\)](#), because the variance of our loss difference is automatically of the same order as its expectation, but the proof is otherwise very similar to theirs, so we omit the details.

LEMMA 3. *Let t_n be a sequence of positive numbers such that $nt_n \rightarrow 0$ and set $G_n = \{\Gamma : R(\Gamma) - R(\Gamma^*) \leq Ct_n\}$ for some $C > 0$. Then $\int \exp[-n\{R_n(\Gamma) - R_n(\Gamma^*)\}] \Pi(d\Gamma) \gtrsim \Pi(G_n) \exp(-2nt_n)$ with P_{Γ^*} -probability converging to 1 as $n \rightarrow \infty$.*

Our last lemma is a summary of various results derived in [Li and Ghosal \(2017\)](#) towards proving their Theorem 3.3. This helps us to fill in the details for the lower bound in Lemma 3 and to identify a sieve with high prior probability but relatively low complexity.

LEMMA 4. Let ε_n be as in Theorem 1 and let $D_n = (n/\log n)^{\frac{1}{\alpha+1}}$. Then, $\|\gamma^* - \hat{\gamma}_{D_n, \beta^*}\|_\infty \leq C\varepsilon_n$ for some $C > 0$, $\beta^* = \beta_{D_n}^* \in (\mathbb{R}^+)^{D_n}$ from (10). Define the following two sets,

$$B_n^* = \{(\beta, d) : \beta \in \mathbb{R}^d, d = D_n, \|\gamma^* - \hat{\gamma}_{d, \beta}\|_\infty \leq C\varepsilon_n\}$$

$$\Sigma_n(r) = \{\gamma : \gamma = \hat{\gamma}_{d, \beta}, \beta \in \mathbb{R}^d, d \leq D_n, \|\beta\|_\infty \leq r\}, \quad r > 0.$$

1. $\Pi(B_n^*) \gtrsim \exp(-an\varepsilon_n)$ for some $a > 0$ depending on C .
2. Set $\Sigma_n = \Sigma_n(K_0 m_n)$ where m_n satisfies $D_n \log m_n \lesssim n\varepsilon_n$ and $K_0 > 0$ is a constant. Then $\Pi(\Sigma_n^c) \lesssim \exp(-Kn\varepsilon_n)$ for some $K > 0$.
3. If $D_n \exp(-Cm_n^t) \lesssim \exp(-n\varepsilon_n)$, for $t > 0$ as in Assumption 3, then the bracketing entropy of Σ_n satisfies $\log N(\varepsilon_n, \Sigma_n, \|\cdot\|_\infty) \lesssim n\varepsilon_n$.

Based on our choice of prior and Assumption 3, $t = 1$, $m_n = n$, and $C > 1$ satisfy the conditions in Lemma 4.

A.4. Proof of Theorem 1.

Define the set

$$(12) \quad A_n = \{\Gamma : \lambda(\Gamma^* \triangle \Gamma) > M\varepsilon_n\}.$$

For the sieve Σ_n in Lemma 4, we have $\Pi_n(A_n) = \Pi_n(A_n \cap \Sigma_n^c) + \Pi_n(A_n \cap \Sigma_n)$. We want to show that both terms in the right hand side of the previous equality vanish, in $L_1(P_{\Gamma^*})$.

It helps to start with a lower bound on $I_n = \int e^{-n\{R_n(\Gamma) - R_n(\Gamma^*)\}} \Pi(d\Gamma)$, the denominator in both of the terms discussed above. First, write

$$I_n \geq \int_{G_n} e^{-n\{R_n(\Gamma) - R_n(\Gamma^*)\}} \Pi(d\Gamma)$$

where G_n is defined in Lemma 3 with $t_n = \varepsilon_n$ and $C > 0$ to be determined. From Proposition 2 and Lemma 1

$$R(\Gamma) - R(\Gamma^*) \leq P_g(X \in \Gamma \triangle \Gamma^*) V \leq \frac{1}{2} V \bar{g} \operatorname{diam}(\Omega) \|\gamma - \gamma^*\|_1$$

where $V = \min\{k - kF_{\Gamma^*}(z) - cF_{\Gamma^{*c}}(z), cF_{\Gamma^{*c}}(z) + kF_{\Gamma^*}(z)\}$. Let $B_\infty(\gamma^*; r)$ denote the set of regions Γ with boundary functions $\gamma = \partial\Gamma$ that satisfy $\|\gamma - \gamma^*\|_\infty \leq r$. If $\Gamma \in B_\infty(\gamma^*; C_0\varepsilon_n)$, then we have

$$\|\gamma - \gamma^*\|_1 \leq 2\pi \bar{g} C_0 \varepsilon_n$$

and, therefore, $R(\Gamma) - R(\Gamma^*) \leq C\varepsilon_n$, where $C = C_0\pi V \bar{g}^2 \operatorname{diam}(\Omega)$. From Lemma 3, we have $I_n \gtrsim \Pi(G_n) e^{-2C\varepsilon_n}$, with P_{Γ^*} -probability converging to 1. Since $G_n \supseteq B_\infty(\gamma^*; C_0\varepsilon_n)$, it follows from Lemma 4, part 1,

$$I_n \gtrsim \Pi\{B_\infty(\gamma^*; C_0\varepsilon_n)\} e^{-2C\varepsilon_n} \gtrsim e^{-C_1 n \varepsilon_n},$$

with P_{Γ^*} probability converging to one, and where $C_1 > 0$ is a constant depending on C_0 and C .

Now we are ready to bound $\Pi_n(A_n \cap \Sigma_n^c)$. Write this quantity as

$$\Pi_n(A_n \cap \Sigma_n^c) = \frac{N_n(A_n \cap \Sigma_n^c)}{I_n} = \frac{1}{I_n} \int_{A_n \cap \Sigma_n^c} e^{-n\{R_n(\Gamma) - R_n(\Gamma^*)\}} \Pi(d\Gamma).$$

It will suffice to bound the expectation of $N_n(A_n \cap \Sigma_n^c)$. By Tonelli's theorem, independence, Lemma 2, and the definition of A_n , we have

$$\begin{aligned} P_{\Gamma^*} N_n(A_n \cap \Sigma_n^c) &= \int_{A_n \cap \Sigma_n^c} P_{\Gamma^*} e^{-n\{R_n(\Gamma) - R_n(\Gamma^*)\}} \Pi(d\Gamma) \\ &= \int_{A_n \cap \Sigma_n^c} \{P_{\Gamma^*} e^{-(\ell_\Gamma - \ell_{\Gamma^*})}\}^n \Pi(d\Gamma) \\ &\leq \int_{A_n \cap \Sigma_n^c} \{1 - \rho \lambda(\Gamma \Delta \Gamma^*)\}^n \Pi(d\Gamma) \\ &\leq \Pi(\Sigma_n^c) e^{-\rho M n \varepsilon_n}. \end{aligned}$$

By Lemma 4, part 2, we have that $\Pi(\Sigma_n^c) \leq e^{-K n \varepsilon_n}$ and, consequently, we can conclude $P_{\Gamma^*} N_n(A_n \cap \Sigma_n^c) \leq e^{-(K + \rho M) n \varepsilon_n}$.

Next, we bound $\Pi_n(A_n \cap \Sigma_n)$ and, as before, it will suffice to bound the expectation of $N_n(A_n \cap \Sigma_n)$. Choose a covering $A_n \cap \Sigma_n$ by sup-norm balls $B_j = B_\infty(\gamma_j; \omega M \varepsilon_n)$, $j = 1, \dots, J_n$, with centers $\gamma_j = \partial \Gamma_j$ in A_n and radii $\omega M \varepsilon_n$, where $\omega < \{\pi \operatorname{diam}(\Omega)\}^{-1}$ is as in Lemma 1. Also, from Lemma 4, part 3, we have that J_n is bounded by $e^{K_1 n \varepsilon_n}$ for some constant $K_1 > 0$. For this covering, we immediately get

$$P_{\Gamma^*} N_n(A_n \cap \Sigma_n) \leq \sum_{j=1}^{J_n} P_{\Gamma^*} N_n(B_j).$$

For each j , using Tonelli, independence, and Lemma 2 again, we find

$$P_{\Gamma^*} N_n(B_j) = \int_{B_j} \{P_{\Gamma^*} e^{-(\ell_\Gamma - \ell_{\Gamma^*})}\}^n \Pi(d\Gamma) \leq \int_{B_j} e^{-n \rho \lambda(\Gamma \Delta \Gamma^*)} \Pi(d\Gamma).$$

By Lemma 1, for Γ in B_j , since the center γ_j is in A_n , it follows that $\lambda(\Gamma \Delta \Gamma^*)$ is lower bounded by $\eta M \varepsilon_n$, where $\eta = \eta(\Gamma^*)$ is given by

$$\eta = \frac{4}{\underline{\gamma}^*} \left(\frac{1}{\operatorname{diam}(\Omega)} - \pi \omega \right) > 0.$$

Therefore, from the bound on J_n ,

$$P_{\Gamma^*} N_n(A_n \cap \Sigma_n) \leq \sum_{j=1}^{J_n} P_{\Gamma^*} N_n(B_j) \leq e^{-n\eta M \varepsilon_n} J_n \leq e^{-(\eta M - K_1)n \varepsilon_n}.$$

Finally, putting everything together, we have

$$\begin{aligned} \Pi_n(A_n) &\leq \frac{N_n(A_n)}{I_n} 1(I_n > e^{-C_1 n \varepsilon_n}) + 1(I_n \leq e^{-C_1 n \varepsilon_n}) \\ &\leq e^{C_1 n \varepsilon_n} \{N_n(A_n \cap \Sigma_n) + N_n(A_n \cap \Sigma_n^c)\} + 1(I_n \leq e^{-C_1 n \varepsilon_n}) \end{aligned}$$

Now take P_{Γ^*} -expectation of both sides:

$$\begin{aligned} P_{\Gamma^*} \Pi_n(A_n) &\leq e^{C_1 n \varepsilon_n} \{P_{\Gamma^*} N_n(A_n \cap \Sigma_n) + P_{\Gamma^*} N_n(A_n \cap \Sigma_n^c)\} + o(1) \\ &\leq e^{-(\eta M - K_1 - C_1)n \varepsilon_n} + e^{-(K + \rho M - C_1)n \varepsilon_n} + o(1). \end{aligned}$$

If $M > \max\{(K_1 + C_1)/\eta, (K - C_1)/\rho\}$, then the upper bound vanishes, completing the proof.

REFERENCES

ANAM, S., UCHINO, E. and SUETAKE, N. (2013). Image boundary detection using the modified level set method and a diffusion filter. *Procedia Comput. Sci.* **22** 192-200.

BISSIRI, P. G., HOLMES, C. C. and WALKER, S. G. (2016). A general framework for updating belief distributions. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 1103–1130. [MR3557191](#)

C., F. M., L., P. E., ADAM, P., JOSEPH, E., A., W. L., P., C. B. and M., E. A. (2010). Ecological boundary detection using Bayesian areal wombling. *Ecology* **91** 3448-3455.

CATONI, O. (1997). *PAC-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning* **56**. Institute of Mathematical Statistics.

DENISON, D. G. T., MALLICK, B. K. and SMITH, A. F. M. (1998). Automatic Bayesian curve fitting. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **60** 333–350. [MR1616029](#)

DiMATTEO, I., GENOVESE, C. R. and KASS, R. E. (2001). Bayesian curve-fitting with free-knot splines. *Biometrika* **88** 1055–1071. [MR1872219](#)

GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732. [MR1380810](#)

GU, K., PATI, D. and DUNSON, D. B. (2014). Bayesian multiscale modeling of closed curves in point clouds. *J. Amer. Statist. Assoc.* **109** 1481–1494. [MR3293605](#)

JIANG, W. and TANNER, M. A. (2008). Gibbs posterior for variable selection in high-dimensional classification and data mining. *Ann. Statist.* **36** 2207–2231. [MR2458185](#)

LI, M. and GHOSAL, S. (2017). Bayesian detection of image boundaries. *Ann. Statist.* **45** 2190–2217. [MR3718166](#)

LI, C., XU, C., GUI, C. and FOX, M. D. (2010). Distance regularized level set evolution and its application to image segmentation. *IEEE Trans. Image Process.* **19** 3243–3254. [MR2789710](#)

LIANG, S., BANERJEE, S. and CARLIN, B. P. (2009). Bayesian wombling for spatial point processes. *Biometrics* **65** 1243–1253. [MR2756512](#)

LU, H. and CARLIN, B. P. (2004). Bayesian areal wombling for geographical boundary analysis. *Geogr. Anal.* **37** 265-285.

MA, H. and CARLIN, B. P. (2007). Bayesian multivariate areal wombling for multiple disease boundary analysis. *Bayesian Anal.* **2** 281–302. [MR2312282](#)

MAINI, R. and AGGARWAL, H. (2009). Study and comparison of various image edge detection techniques. *IJIP* **3** 1-11.

MAMMEN, E. and TSYBAKOV, A. B. (1995). Asymptotical minimax recovery of sets with smooth boundaries. *Ann. Statist.* **23** 502–524. [MR1332579](#)

MARTIN, D. R., FOWLKES, C. C. and MALIK, J. (2004). Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. Pattern Anal. Mach. Intell.* **26** 530-549.

SCHUMAKER, L. L. (2007). *Spline Functions Basic Theory*. Wiley, New York.

SHEN, W. and GHOSAL, S. (2015). Adaptive Bayesian procedures using random series priors. *Scand. J. Stat.* **42** 1194–1213. [MR3426318](#)

SHEN, X. and WASSERMAN, L. (2001). Rates of convergence of posterior distributions. *Ann. Statist.* **29** 687–714. [MR1865337](#)

SYRING, N. and LI, M. (2017a). BayesBD: Bayesian inference for image boundaries R package version 1.2.

SYRING, N. and LI, M. (2017b). BayesBD: an R package for Bayesian inference on image boundaries. *The R Journal* **9** 149-162.

SYRING, N. and MARTIN, R. (2017). Gibbs posterior inference on the minimum clinically important difference. *J. Statist. Plann. Inference* **187** 67–77. [MR3638043](#)

SYRING, N. and MARTIN, R. (2018). Calibrating posterior credible regions. *Biometrika*, to appear.

YUAN, W., CHIN, K. S., HUA, M., DONG, G. and WANG, C. (2016). Shape classification of wear particles by image boundary analysis using machine learning algorithms. *Mech. Syst. Signal Pr.* **72-73** 346-358.

ZHANG, T. (2006). Information-theoretic upper and lower bounds for statistical estimation. *IEEE T. Inform. Theory* **52** 1307-1321.

ZIOU, D. and TABBONE, S. (1998). Edge detection techniques – an overview. *Pattern Recognition and Image Analysis C/C of Raspoznavaniye Obrazov I Analiz Izobrazhenii* **8** 537-559.

DEPARTMENT OF MATHEMATICS AND STATISTICS
 WASHINGTON UNIVERSITY IN ST. LOUIS
 CUPPLES 1
 ONE BROOKINGS DRIVE
 ST. LOUIS, MO 63130-4899
 E-MAIL: nasyring@wustl.edu

DEPARTMENT OF STATISTICS
 NORTH CAROLINA STATE UNIVERSITY
 5109 SAS HALL
 2311 STINSON DR.
 RALEIGH, NC 27695
 E-MAIL: rgmarti3@ncsu.edu