A MONTE CARLO GAP TEST IN COMPUTING HPD REGIONS

MING-HUI CHEN*
Department of Statistics, University of Connecticut, 215 Glenbrook Road, Storrs, CT 06269, USA
E-mail: mhchen@stat.uconn.edu

XUMING HE
Department of Statistics, University of Illinois, 725 S. Wright, Champaign, Illinois 61820, USA
E-mail: x-he@uiuc.edu

QI-MAN SHAO
Department of Mathematics, University of Oregon, Eugene, OR 97403-1222, USA
E-mail: qmshao@darkwing.uoregon.edu

HAI XU
Department of Statistics, University of Connecticut, 215 Glenbrook Road, Storrs, CT 06269, USA
E-mail: haixu@stat.uconn.edu

We consider estimation of the Bayesian highest posterior density (HPD) regions for a parameter of interest and propose a Monte Carlo gap test based on a random or dependent sample of the relevant parameter generated from its posterior distribution to determine whether the Bayesian regions contain one or more intervals. We also provide a simple Monte Carlo procedure to determine the exact form of the HPD region based on the outcome of the gap test. The basic theory is developed and a simulation study is conducted in examining the performance of the Monte Carlo gap test.

1 Introduction

Highest posterior density (HPD) region is a useful summary of a posterior distribution for some parameter of interest. Suppose the parameter is one-dimensional. As discussed in Box and Tiao (1992), an HPD region has two main properties:

(a) the density for every point inside the region is greater than that for every point outside the region; and

*CONTRIBUTED TO "MODERN STATISTICS (ON THE OCCASION OF PROF. YIAO-TING ZHANG'S 70TH BIRTHDAY)", EDS. JIAN HUANG AND HEPING ZHANG

(b) for a given probability content, say $1-\alpha$, the region is of the smallest size.

Unlike other Bayesian posterior quantities, such as the posterior mean and the posterior variance, it is typically computationally intensive to compute the HPD region unless one is dealing with a simple model such as the standard normal model.

Consider a Bayesian posterior density of the form

$$\pi(\theta, \varphi|D) = \frac{1}{c(D)} L(\theta, \varphi|D)\pi(\theta, \varphi), \tag{1}$$

where D denotes *data*, the parameter θ is one-dimensional, and φ may be a multidimensional vector of parameters other than θ in the model. In (1), $L(\theta, \varphi|D)$ is a likelihood function given D, $\pi(\theta, \varphi)$ is a prior, and $c(D)$ is the normalizing constant. Our objective is to obtain a Bayesian HPD region for θ.

Let $\pi(\theta|D)$ and $\Pi(\theta|D)$ denote the marginal posterior density function and the marginal posterior cumulative distribution function (cdf) of θ, respectively. One of the widely used $100(1-\alpha)\%$ Bayesian credible intervals for θ takes the form

$$(\theta^{(\alpha/2)}, \theta^{(1-\alpha/2)}),$$

where

$$\Pi(\theta^{(\alpha/2)}|D) = \alpha/2 \text{ and } \Pi(\theta^{(1-\alpha/2)}|D) = 1-\alpha/2. \tag{2}$$

When $\pi(\theta|D)$ is symmetric and unimodal, the Bayesian credible interval $(\theta^{(\alpha/2)}, \theta^{(1-\alpha/2)})$ is also an HPD interval. However, when $\pi(\theta|D)$ is not symmetric, $(\theta^{(\alpha/2)}, \theta^{(1-\alpha/2)})$ is not an HPD interval in general. In this case, an HPD interval or region is more desirable, as it better displays the important features of the posterior distribution than does a credible interval. A $100(1-\alpha)\%$ HPD region for θ is given by

$$R(\pi_\alpha) = \{\theta : \pi(\theta|D) \geq \pi_\alpha\}, \tag{3}$$

where π_α is the largest constant such that $P(\theta \in R(\pi_\alpha)) \geq 1-\alpha$. In (3), $R(\pi_\alpha)$ can be reduced to a $100(1-\alpha)\%$ HPD interval for θ when $\pi(\theta|D)$ is unimodal.

Due to the recent advances in computing technology and the development of Markov chain Monte Carlo (MCMC) sampling algorithms, several Monte Carlo algorithms have been proposed for computing Bayesian HPD intervals or regions. In particular, Tanner (1996), Wei and Tanner (1990), and Hyndman (1996) provide a Monte Carlo (MC) algorithm to calculate the content and boundary of HPD regions. However, their MC algorithm requires evaluating the marginal posterior densities analytically or numerically. The implementation of their algorithm is also quite complicated and computationally intensive. Chen and Shao (1999) propose a simpler MC method for computing HPD intervals. Their method does not require knowing the closed forms of $\pi(\theta|D)$ and $\Pi(\theta|D)$, and can be applied to compute HPD intervals for the parameters of interest and also for functions of the parameters. However, the algorithm proposed by Chen and Shao (1999) requires the unimodality of the marginal posterior distribution $\pi(\theta|D)$.

We note that the marginal posterior density $\pi(\theta|D)$ is given by

$$\pi(\theta|D) = \int \pi(\theta, \boldsymbol{\varphi}|D) \, d\boldsymbol{\varphi} = \int \frac{1}{c(D)} L(\theta, \boldsymbol{\varphi}|D) \pi(\theta, \boldsymbol{\varphi}) \, d\boldsymbol{\varphi}, \tag{4}$$

where the normalizing constant $c(D)$ is often unknown and the analytical evaluation of the integral may not be available. This is particularly true when $\boldsymbol{\varphi}$ is high-dimensional. Therefore, when $\pi(\theta|D)$ is analytically intractable, it is often difficult to know whether $\pi(\theta|D)$ is unimodal or multimodal.

We also note that when $\pi(\theta|D)$ is multimodal, the HPD region often consists of several intervals. In such cases, there exist gaps between the intervals inside the HPD region. A challenging problem is how to determine whether there are any gaps inside the HPD region based on a Monte Carlo sample from $\pi(\theta|D)$. In this article, a novel Monte Carlo gap test is proposed for determining possible gaps inside an HPD region. As a byproduct, a Monte Carlo algorithm is also developed for determining the end points of intervals inside the HPD region.

The remainder of the paper is organized as follows. In Section 2, we give a brief overview of the current Monte Carlo methods for computing HPD intervals/regions. The main results of the Monte Carlo gap test are presented in Section 3. The performance of the Monte Carlo gap test is examined in detail via a simulation study in Section 4. Some concluding remarks are given in Section 5.

2 Current Monte Carlo Methods

2.1 Density Quantile Approach

This approach was first introduced by Wei and Tanner (1990) and further elaborated by Hyndman (1996). Let $\{(\theta_i, \boldsymbol{\varphi}_i), \quad i = 1, 2, \ldots, n\}$ denote a random (or MCMC) sample from $\pi(\theta, \boldsymbol{\varphi}|D)$. Then, $\{\theta_i, \quad i = 1, 2, \ldots, n\}$ is a random (or MCMC) sample from $\pi(\theta|D)$. Also, let $\xi_i = \pi(\theta_i|D)$ for

$i = 1, 2, \ldots, n$, and let $\xi_{([\alpha n])}$ be the $[\alpha n]^{\text{th}}$ smallest of $\{\xi_i\}$.

Define

$$R_n(\alpha) = \{\theta : \ \pi(\theta|D) \geq \xi_{([\alpha n])}\}, \tag{5}$$

which is the set of θ's such that their density values are greater than or equal to $\xi_{([\alpha n])}$. In Hyndman (1996), it was stated that $R_n(\alpha)$ converges to $R(\pi_\alpha)$ given in (3), as $n \to \infty$ when $\{\theta_i, \ i = 1, 2, \ldots, n\}$ is a random sample from $\pi(\theta|D)$. However, no formal proof was provided by Hyndman (1996). Wei and Tanner (1990) proposed a Monte Carlo method for calculating the boundary of an HPD region via data augmentation. Hyndman (1996) discussed how to compute $R_n(\alpha)$ using the contour function for a bivariate HPD region and the grid method with spline interpolation for a univariate HPD region.

2.2 Sampling Quantile Approach

Assume that $\pi(\theta|D)$ is unimodal and $\{(\theta_i, \varphi_i), \ i = 1, 2, \ldots, n\}$ is an MCMC sample from the joint posterior distribution $\pi(\theta, \varphi|D)$.

As discussed earlier, a $100(1-\alpha)\%$ HPD interval must have a probability content of $1 - \alpha$ and the density for every point inside the interval must be greater than that for every point outside the interval. Therefore, an HPD interval is a special credible interval and this interval is of the shortest length among all possible credible intervals with the same probability content $1 - \alpha$. Based on the main properties of the HPD interval, Chen and Shao (1999) proposed the following procedure for calculating an HPD interval for θ:

Chen–Shao HPD Estimation Algorithm

Step 1. Obtain an MCMC sample $\{\theta_i, \ i = 1, 2, \ldots, n\}$ from $\pi(\theta|D)$.

Step 2. Sort $\{\theta_i, \ i = 1, 2, \ldots, n\}$ to obtain the ordered values:

$$\theta_{(1)} \leq \theta_{(2)} \leq \cdots \leq \theta_{(n)}.$$

Step 3. Compute the $100(1-\alpha)\%$ credible intervals

$$R_j(n) = (\theta_{(j)}, \theta_{(j+[(1-\alpha)n])})$$

for $j = 1, 2, \ldots, n - [(1-\alpha)n]$.

Step 4. The $100(1-\alpha)\%$ HPD interval is the one, denoted by $R_{j^*}(n)$, with *the smallest interval width* among all credible intervals.

Under certain regularity conditions, Chen and Shao (1999) showed that the above procedure is asymptotically valid. We state this result in the following proposition.

Proposition 2.1 *Assume that* $\{\theta_i,\ i = 1, 2, \ldots, n\}$ *is an ergodic MCMC sample from* $\pi(\theta|D)$. *If* $\pi(\theta|D)$ *is unimodal and*

$$\min_{\theta_L < \theta_U} \left(|\pi(\theta_U|D) - \pi(\theta_L|D)| + |\Pi(\theta_U|D) - \Pi(\theta_L|D) - (1-\alpha)| \right)$$

has a unique solution, then

$$R_{j^*}(n) \to R(\pi_\alpha) \ a.s. \ \text{ as } n \to \infty,$$

where $R(\pi_\alpha)$ *is defined in* (3).

The proof of this proposition can be found in Chen and Shao (1999). Chen and Shao (1999) also considered the extension of their method to the multimodal case. The following conjecture is given in Chen and Shao (1999).

Conjecture 2.1 *Assume that* $\pi(\theta|D)$ *has at most two modes. Let* $\{\theta_1, \theta_2, \ldots, \theta_n\}$ *be a sample from* $\pi(\theta|D)$ *and let* $\theta_{(j)}$ *be the ordered values of the* θ_j. *For* $0 < \alpha < 1$, *denote*

$$\begin{aligned} D = & \min_{0 \le m \le [(1-\alpha)n]} \ \min_{0 \le i \le n - [n(1-\alpha)] - 2m} \\ & \times \min_{i+m \le j \le n - [n(1-\alpha)] - m} \{(\theta_{(i+m)} - \theta_{(i)}) + (\theta_{(j+[n\alpha]-m)} - \theta_{(j)})\}. \end{aligned}$$

If $(\theta_{(i^*+m^*)} - \theta_{(i^*)}) + (\theta_{(j^*+[n\alpha]-m^*)} - \theta_{(j^*)}) = D$, *then* $(\theta_{(i^*)},\ \theta_{(i^*+m^*)}) \cup (\theta_{(j^*)},\ \theta_{(j^*+[n\alpha]-m^*)})$ *is an approximate HPD region for* θ.

It is expected that when $\pi(\theta|D)$ is unimodal, $(\theta_{(i^*)},\ \theta_{(i^*+m^*)}) \cup (\theta_{(j^*)},\ \theta_{(j^*+[n\alpha]-m^*)})$ can be automatically reduced to one interval. Similarly, this conjecture can be extended to cases where $\pi(\theta|D)$ has more than two modes. However, the question remains open as to how one can test the unimodality of $\pi(\theta|D)$ based on an MC sample $\{\theta_i,\ i = 1, 2, \ldots, n\}$. We will address this issue in the next section.

3 Monte Carlo Gap Tests

To develop a Monte Carlo gap test, we first show that the density quantile estimate of $R_n(\alpha)$ given in (5) is asymptotically correct. That is, $R_n(\alpha) \to$

$R(\pi_\alpha)$, as $n \to \infty$, where $R(\pi_\alpha)$ is defined in (3). We formally state this result in the following theorem.

Theorem 3.1 *Let $\pi_n = \xi_{([\alpha n])}$ denote the lower α^{th} sample quantile of $\xi_i = \pi(\theta_i|D)$. Then,*

$$\pi_n \to \pi_\alpha \quad as \quad n \to \infty.$$

Consequently, $R_n(\alpha) \to R(\pi_\alpha)$, as $n \to \infty$.

Proof: First note that $\pi_\alpha > 0$ for any $\alpha > 0$. Since π_n is the sample quantile of $\{\pi(\theta_i)\}$ and π_α is the corresponding population quantile, the theorem follows from the standard quantile convergence result. $\square$

Let

$$\Omega_n(\alpha) = \{\theta_i : \quad \xi_i = \pi(\theta_i|D) \geq \xi_{([\alpha n])}\}, \tag{6}$$

be the set of θ_i's such that their density values are greater than or equal to $\xi_{([\alpha n])}$. Denote

$$\theta_\alpha^L(n) = \min_{\{\theta_i \in \Omega_n(\alpha)\}} \{\theta_i\}, \quad \text{and} \quad \theta_\alpha^U(n) = \max_{\{\theta_i \in \Omega_n(\alpha)\}} \{\theta_i\}. \tag{7}$$

REMARK 3.1: If $\pi(\theta|D)$ is unimodal, then $R(\pi_\alpha)$ reduces to an interval, called the HPD interval. In this case, we have

$$R_\alpha(n) = (\theta_\alpha^L(n),\ \theta_\alpha^U(n)), \tag{8}$$

and by Theorem 3.1, $(\theta_\alpha^L(n),\ \theta_\alpha^U(n))$ is a consistent estimator of the HPD interval $R(\pi_\alpha)$.

Definition 3.1: Let $R(\pi_\alpha)$ denote a $100(1-\alpha)\%$ HPD region. There are no gaps in $R(\pi_\alpha)$ if there exist constants $a < b$ such that $P(R(\pi_\alpha) = (a,b)|D) = 1$.

Roughly speaking, Definition 3.1 implies that if there are no gaps, $R(\pi_\alpha)$ reduces to an interval. On the other hand, if $R(\pi_\alpha)$ consists of several disconnected intervals, then there are gaps. For example, if $\pi(\theta|D)$ is bimodal, typically $R(\pi_\alpha)$ consists of two intervals, and therefore there is a gap inside $R(\pi_\alpha)$. But, we notice that it is not necessary that a bimodal distribution always leads to a gap inside $R(\pi_\alpha)$. Whether there is a gap in $R(\pi_\alpha)$ depends on the shape of the density and the choice of $1-\alpha$. For example, if the density values are similar at the two modes, then the larger $1-\alpha$ or the shorter the distance between the two modes, the less likely there is a gap.

Next, we introduce a useful lemma.

Lemma 3.1 *Let $\{X_i, 1 \leq i \leq n\}$ be a random sample of size n from a population distribution F. Suppose that the density function f of X_i exists and is continuous on R^1 with $f(x) > 0$ for x in the support of X_i. Assume that for any $a < b$, there exist $C > 0, r > 0$ such that $|f(x) - f(y)| \leq C|x-y|^r$ for any $a \leq x < y \leq b$. Let $X_{n,1} \leq X_{n,2} \leq \cdots \leq X_{n,n}$ be the order statistics. Then for any $0 < \alpha_1 < \alpha_2 < 1$*

$$P(n \max_{\alpha_1 n \leq i \leq \alpha_2 n} (X_{n,i} - X_{n,i-1}) f(X_{n,i}) \leq \log n + x) \to e^{-(\alpha_2 - \alpha_1) e^{-x}}. \tag{9}$$

Proof: The outline of the proof is given as follows. Let $U_i = F(X_i)$ and let $\{\xi_i, 1 \leq i \leq n+1\}$ be i.i.d. exponential random variables with mean 1. Put $S_k = \sum_{i=1}^k \xi_i$. Then U_i are i.i.d. uniformly distributed random variables over $(0,1)$ and $\{U_{n,i}, 1 \leq i \leq n\}$ and $\{S_i/S_{n+1}, 1 \leq i \leq n\}$ have the same distribution. Observe that

$$F(X_{n,i}) - F(X_{n,i-1}) = f(X_{n,i} - \delta_i)(X_{n,i} - X_{n,i-1}), \tag{10}$$

where $0 \leq \delta_i \leq X_{n,i} - X_{n,i-1}$. It is easy to see that for any $0 < \alpha_1 < \alpha_2 < 1$, there exist $a_1 < a_2$ such that

$$a_1 \leq X_{n,[\alpha_1 n]} \leq X_{n,[\alpha_2 n]} \leq a_2 \quad a.s.$$

as $n \to \infty$. Thus, by (10) and by the assumption that f is continuous and positive in the support of X_i

$$\begin{aligned} &\max_{\alpha_1 n \leq i \leq \alpha_2 n} (X_{n,i} - X_{n,i-1}) \\ &= O(1) \max_{\alpha_1 n \leq i \leq \alpha_2 n} (F(X_{n,i}) - F(X_{n,i-1})) = O(n^{-1} \log n) \quad a.s. \end{aligned}$$

Therefore, by the Lipschitz condition for f

$$\begin{aligned} &n \max_{\alpha_1 n \leq i \leq \alpha_2 n} (X_{n,i} - X_{n,i-1}) f(X_{n,i}) - n \max_{\alpha_1 n \leq i \leq \alpha_2 n} (U_{n,i} - U_{n,i-1}) \qquad (11) \\ &= O(n) \max_{\alpha_1 n \leq i \leq \alpha_2 n} (X_{n,i} - X_{n,i-1})^{1+r} = O(n^{-r} (\log n)^{1+r}) = o(1) \quad a.s. \end{aligned}$$

It is easy to show that $n/S_{n+1} \to 1$, and

$$n \max_{\alpha_1 n \leq i \leq \alpha_2 n} \xi_i / S_{n+1} - \max_{\alpha_1 n \leq i \leq \alpha_2 n} \xi_i \to 0. \tag{12}$$

Furthermore,

$$\begin{aligned} &P\Big(\max_{\alpha_1 n \leq i \leq \alpha_2 n} \xi_i \leq \log n + x\Big) = (1 - e^{-(\log n + x)})^{(\alpha_2 - \alpha_1) n} \\ &= \Big(1 - e^{-x}/n\Big)^{(\alpha_2 - \alpha_1) n} \to \exp(-(\alpha_2 - \alpha_1) e^{-x}) \end{aligned}$$

as $n \to \infty$. The lemma then follows from (11), (12), and the fact that $\{U_{n,i}, 1 \le i \le n\}$ and $\{S_i/S_{n+1}, 1 \le i \le n\}$ have the same distribution. $\square$

Let $a_n = \alpha_{1,n}n$ denote the number of θ_i's that are less than equal to $\theta_\alpha^L(n)$ and $b_n = \alpha_{2,n}n$ denote the number of θ_i's that are less than equal to $\theta_\alpha^U(n)$, where $\theta_\alpha^L(n)$ and $\theta_\alpha^U(n)$ are defined in (7). We are led to the following theorem.

Theorem 3.2 *Assume that $\pi(\theta|D)$ is continuous and $R(\pi_\alpha)$ has no gaps. Also assume that for any $a < b$, there exist $C > 0, r > 0$ such that $|\pi(\theta|D) - \pi(\theta^*|D)| \le C|\theta - \theta^*|^r$ for any $a \le \theta < \theta^* \le b$. Suppose $\{\theta_i,\ i = 1, 2, \ldots, n\}$ is a random sample from $\pi(\theta|D)$. Let $\theta_{n,1} \le \theta_{n,2} \le \cdots \le \theta_{n,n}$ be the order statistics. Then,*

$$P(n \max_{a_n < i \le b_n} (\theta_{n,i} - \theta_{n,i-1})\pi(\theta_{n,i}|D) \le \log n + x) \to e^{-(1-\alpha)e^{-x}}. \qquad (13)$$

Proof: Notice that when $R(\pi_\alpha)$ has no gaps, by the definition of $\Omega_n(\alpha)$, $\theta_{n,i} \in \Omega_n(\alpha)$ for all i's such that $a_n \le i \le b_n$. Also, it is easy to show that $\alpha_{2,n} - \alpha_{1,n} = (b_n - a_n)/n \to 1 - \alpha$ as $n \to \infty$. The rest of the proof follows from Lemma 3.1. $\square$

Determining whether there are any gaps inside a $100(1-\alpha)\%$ HPD region $R(\pi_\alpha)$ is a standard hypothesis problem. That is, we wish to test a null hypothesis H_0: no gaps in $R(\pi_\alpha)$ versus an alternative hypothesis H_a: one or more gaps in $R(\pi_\alpha)$. Theorem 3.2 directly leads to the following Monte Carlo gap test.

Monte Carlo Gap Test:

Step 1: Generate a random sample from $\{\theta_i,\ i = 1, 2, \ldots, n\}$ from $\pi(\theta|D)$.

Step 2: Compute $\xi_i = \pi(\theta_i|D)$ and the α^{th} lower sample quantile $\xi_{([\alpha n])}$.

Step 3: Sort $\{\theta_i :\ \theta_i \in \Omega_n(\alpha)\}$, where $\Omega_n(\alpha) = \{\theta_i :\ \pi(\theta_i|D) \ge \xi_{([\alpha n])}\}$, to obtain the ordered values denoted by

$$\theta_{n_\alpha,1} \le \theta_{n_\alpha,2} \le \cdots \le \theta_{n_\alpha,n_\alpha},$$

where n_α is the size of $\Omega_n(\alpha)$.

Step 4. Compute the test statistic

$$T = n \max_{1 < i \le n_\alpha} (\theta_{n_\alpha,i} - \theta_{n_\alpha,i-1})\pi(\theta_{n_\alpha,i}|D) - \log n. \qquad (14)$$

Step 5. Compute the asymptotic p-value

$$p\text{-value} = P(T > t^*) = 1 - e^{-(1-\alpha)e^{-t^*}}, \tag{15}$$

where t^* denotes the observed value of T given in (14).

Step 6. If p-value $\leq \alpha^*$, a prespecified level of significance, we reject H_0, and hence, we conclude that $R(\pi_\alpha)$ has one or more gaps. If p-value $> \alpha^*$, we do not reject H_0, and therefore, we conclude that there are no gaps in $R(\pi_\alpha)$.

REMARK 3.2: Suppose that the gap test leads to the rejection of H_0 and assume i_α is the integer such that

$$(\theta_{n_\alpha,i_\alpha} - \theta_{n_\alpha,i_\alpha-1})\pi(\theta_{n_\alpha,i_\alpha}|D) = \max_{1<i\leq n_\alpha}(\theta_{n_\alpha,i} - \theta_{n_\alpha,i-1})\pi(\theta_{n_\alpha,i}|D).$$

Then, the lower limit of the gap is $\theta_{n_\alpha,i_\alpha-1}$ and the upper limit of the gap is $\theta_{n_\alpha,i_\alpha}$. In this case, we can further apply the gap test for

$$\theta_{n_\alpha,1} \leq \theta_{n_\alpha,2} \leq \cdots \leq \theta_{n_\alpha,i_\alpha-1},$$

and

$$\theta_{n_\alpha,i_\alpha} \leq \theta_{n_\alpha,2} \leq \cdots \leq \theta_{n_\alpha,n_\alpha},$$

respectively. Let t_1^* and t_2^* denote the respective observed values. Then, the p-values can be approximated by

$$1 - e^{-((i_\alpha-1)/n)e^{-t_1^*}} \text{ and } 1 - e^{-((n_\alpha-i_\alpha+1)/n)e^{-t_2^*}},$$

respectively. We continue this process until H_0 is accepted. We notice that a large sample size is required when the number of gaps is large. We also caution making a large number of such tests without controlling for the false positive rate.

REMARK 3.3: Suppose $\pi(\theta|D)$ is bimodal and $R(\pi_\alpha)$ has exactly one gap. Assume that the gap test in fact detects the gap. Using the same notation introduced in Remark 3.2, as a byproduct, $R(\pi_\alpha)$ can be approximated by

$$(\theta_{n_\alpha,1},\ \ \theta_{n_\alpha,i_\alpha-1}) \cup (\theta_{n_\alpha,i_\alpha},\ \ \theta_{n_\alpha,n_\alpha}).$$

REMARK 3.4: In practice, $\pi(\theta|D)$ is unknown. In this case, we propose to use a Monte Carlo estimate $\hat{\pi}(\theta|D)$ to replace $\pi(\theta|D)$ in the gap test.

There are several density estimation methods available in the literature, including the kernel density estimate, the conditional marginal density estimate (CMDE) (see, for example, Gelfand, Smith and Lee (1992)), and the importance-weighted marginal posterior density estimate (IWMDE) of Chen (1994). Let $\{\theta_i, \ \ i=1,2,\dots\}$ denote a random sample from $\pi(\theta|D)$. Then, the kernel density estimate has the form

$$\hat{\pi}_{kernel}(\theta|D) = \frac{1}{nh_n}\sum_{i=1}^{n}\mathcal{K}\left(\frac{\theta-\theta_i}{h_n}\right), \tag{16}$$

where the kernel $\mathcal{K}$ is a bounded density on R^1, h_n is the bandwidth, and θ is a point in R^1. As recommended by Silverman (1986), if a Gaussian kernel, i.e., $\mathcal{K}(\theta)=(1/\sqrt{2\pi})e^{-\theta^2/2}$, is used, a good choice of h_n is $1.06\sigma^* n^{-1/5}$, where σ^* is the sample standard deviation of the θ_i's.

Let $\{(\theta_i,\boldsymbol{\varphi}_i), \quad i=1,2,\dots,n\}$ denote a random or MCMC sample from $\pi(\theta,\boldsymbol{\varphi}|D)$. Then, the conditional marginal density estimate (CMDE) of $\pi(\theta|D)$ has the form

$$\hat{\pi}_{cmde}(\theta|D) = \frac{1}{n}\sum_{i=1}^{n}\pi\left(\theta|\boldsymbol{\varphi}_i, D\right), \tag{17}$$

where $\pi(\theta|\boldsymbol{\varphi},D)$ denotes the conditional posterior density of θ. We note that in (17), $\pi(\theta|\boldsymbol{\varphi},D)$ needs to be completely known. Here, "*completely* known" means that $\pi(\theta|\boldsymbol{\varphi},D)$ can be evaluated at any point of $(\theta,\boldsymbol{\varphi})$. In other words, the kernel *and* the normalizing constant of this conditional density are available in closed form.

An extension to the CMDE given by Chen (1994) is the importance-weighted marginal posterior density estimate (IWMDE). The IWMDE of $\pi(\theta|D)$ takes the following form:

$$\hat{\pi}_{iwmde}(\theta|D) = \frac{1}{n}\sum_{i=1}^{n} w(\theta_i|\boldsymbol{\varphi}_i)\frac{L(\theta,\boldsymbol{\varphi}_i|D)\pi(\theta,\boldsymbol{\varphi}_i)}{L(\theta_i,\boldsymbol{\varphi}_i|D)\pi(\theta_i,\boldsymbol{\varphi}_i)}, \tag{18}$$

where $w(\theta|\boldsymbol{\varphi})$ is a completely known conditional density whose support is contained in or equal to the support of the true conditional density $\pi(\theta|\boldsymbol{\varphi},D)$. See Chen, Shao, and Ibrahim (2000) for more detailed discussions on the density estimation based on Monet Carlo samples.

We will study the performance of the gap test with known and unknown $\pi(\theta|D)$ in detail in the next section.

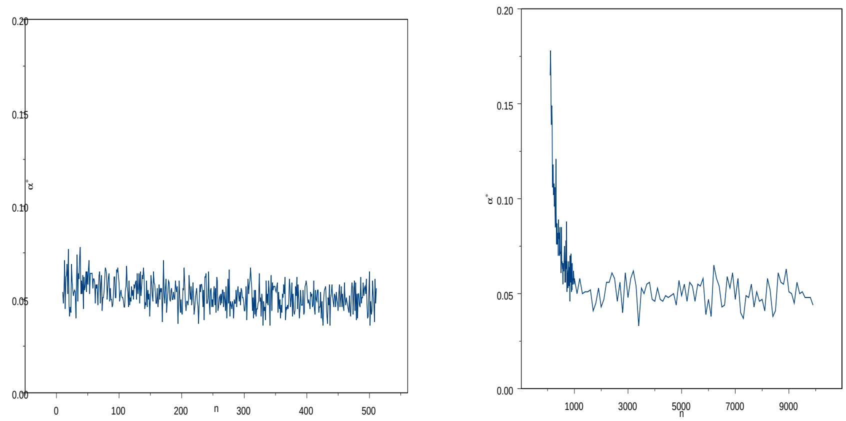

Figure 1. Estimates of Type I Error Probabilities (α^*) under the Bimodal Normal Distribution (left) and under the Bimodal Cauchy (right) for various sample sizes.

4 A Simulation Study

In this section, we conduct an extensive simulation study to examine the performance of the Monte Carlo gap test developed in the previous section. In all simulations, we use $1 - \alpha = 0.95$ for the HPD region.

First we study how large the sample size is required to achieve a prespecified type I error probability α^*. We take $\alpha^* = 0.05$ and we consider $\pi(\theta|D)$ to be a standard normal ($N(0,1)$) density and a standard Cauchy (Cauchy$(0,1)$) density. We generate a random sample of size n from each of these two distributions. Then, we perform the Monte Carlo gap test. Let $I_k = 1$ if the test leads to a rejection and $I_k = 0$ otherwise. We repeat this procedure $K = 1,000$ times so that we obtain $\{I_k, \quad k = 1, 2, \ldots, 1000\}$ and hence, an estimate of α^* is given by $\hat{\alpha}^* = \frac{1}{1000}\sum_{k=1}^{1000} I_k$. The results are displayed in Figure 1. It shows that for the standard normal case, $\hat{\alpha}^*$'s are almost the same as the true α^* for n as small as 50, but for the standard Cauchy case, n is required to be 1000 in order to achieve the true level of the type I error.

Second, we examine the power of the Monte Carlo gap test. Again, we choose $\alpha^* = 0.05$. We consider a bimodal normal $(0.5[N(-\mu/2, 1) + N(\mu/2, 1)])$ density for $\mu \in [3.5, 5]$ and a bimodal Cauchy $(0.5[\text{Cauchy}(-\mu/2, 1) + \text{Cauchy}(\mu/2, 1)])$ density for $\mu \in [20, 40]$. The sample sizes are $n = 1000$ and $n = 5000$ for both cases. We repeat the simulation $K = 1000$ times. The estimated power functions (β) versus the gap size are

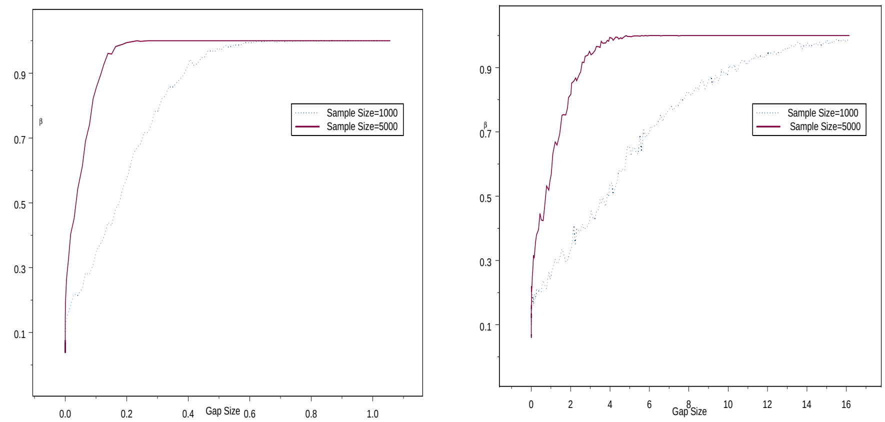

Figure 2. Estimates of the Powers (β) of the Gap Test under the Bimodal Normal Distribution (left) and under the Bimodal Cauchy (right).

Table 1. True and Estimated End-Points

End-Point	True Value	Estimated Value		
		Mean	Std Dev	IQR
1	-3.854	-3.851	0.028	0.038
2	-0.246	-0.249	0.027	0.038
3	0.960	0.964	0.013	0.018
4	3.127	3.123	0.012	0.016

plotted in Figure 2. It clearly demonstrates that for the bimodal normal density, the gap test can detect a very small gap even when $n = 1000$. However, for the Cauchy case, the power curve approaches 1 when the gap size is approximately 4 and 16 for $n = 5000$ and $n = 1000$, respectively. Thus, the gap test performs better for light tailed distributions.

Third, we investigate how accurate our HPD region estimate is based on the gap test. We choose $\pi(\theta|D)$ to be the bimodal normal $0.5 * [N(-2.05, 1) + N(2.05, 0.25)]$ density. The Monet Carlo sample size is $n = 5000$. We repeat the calculations 1000 times. In this case, the 95% HPD region contains two intervals. Therefore, we have four end-points. The true density and the estimates of these four end-points are plotted in Figure 3. A numerical summary of the boxplots shown in Figure 3 is given in Table 1. From both Figure 3 and Table 1, it can be seen that the Monte Carlo estimates are quite good.

Finally, we examine the performance of the Monte Carlo gap test when

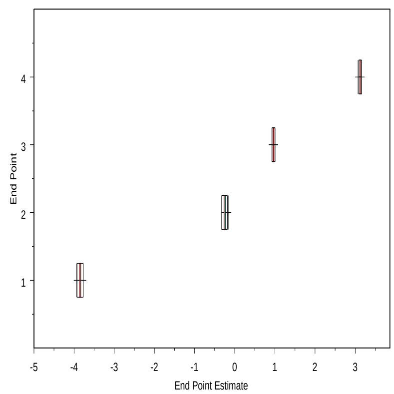

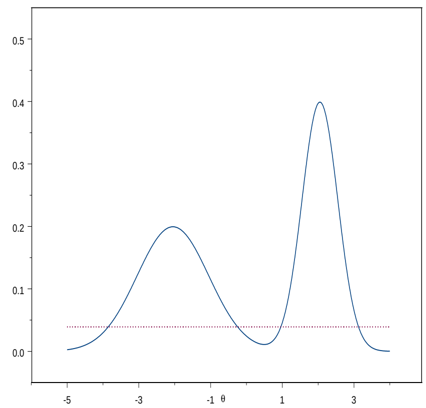

Figure 3. Boxplots of End-Point Estimates of HPD Regions (left) and the True Bimodal Probability Density Function (right).

$\pi(\theta|D)$ is replaced by the kernel density estimate and CMDE. We consider $\pi(\theta, \varphi|D)$ to be a mixture of two bivariate normal density:

$$0.5 * \left[N_2 \left(\begin{pmatrix} -2.1 \\ -2.1 \end{pmatrix}, \Sigma_1 \right) + N_2 \left(\begin{pmatrix} 2.1 \\ 2.1 \end{pmatrix}, \Sigma_2 \right) \right],$$

where $\Sigma_1 = \Sigma_2 = \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix}$. Straightforward algebra shows that the marginal distribution for θ is $0.5[N(-2.1, 1) + N(2.1, 1)]$. For sample size $n = 1000$, the powers calculated based on the true density $\pi(\theta|D)$, the kernel density estimate, and the CMDE for $\pi(\theta|D)$ are $\beta = 0.54$, $\beta_{kernel} = 0.198$, and $\beta_{cmde} = 0.571$, respectively. For sample size $n = 5000$, the powers calculated based on the true density $\pi(\theta|D)$, the kernel density estimate, and the CMDE for $\pi(\theta|D)$ are $\beta = 0.99$, $\beta_{kernel} = 0.542$, and $\beta_{cmde} = 0.984$, respectively. In addition, we take $\pi(\theta, \varphi|D)$ to be $N_2 \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_1 \right)$ so that the marginal distribution for θ is $N(0, 1)$. In this case, the estimated type I error probabilities are $\hat{\alpha}^*_{cmde} = 0.05$ when $n = 50$ for the CMDE and $\hat{\alpha}^*_{kernel} = 0.033$ when $n = 500$ for the kernel density estimate. This simulation study clearly demonstrates that with the CMDE, the performance of the gap test with an estimated density is as good as the one with the known density. However, the gap test with the kernel density estimate performs less favorably.

5 Concluding Remarks

In this paper, we proposed a novel Monte Carlo gap test. The aim of the gap test is to determine whether an HPD region consists of only one or more intervals using a Monte Carlo sample from the posterior distribution. The simulation study presented in Section 4 shows that the gap test performs well with modestly large sample sizes. Today's computer power has made it relatively easy to generate thousands of random numbers from a posterior distribution, so the sample size requirement for the proposed gap test is quite realistic.

Although we only consider a random sample from $\pi(\theta|D)$, our results may be extended a dependent MCMC sample. However, the performance of the gap test with dependent samples needs further studies. Extension of the gap test to higher dimensional parameters will also be worthwhile.

References

1. G.E.P. Box and G.C. Tiao, *Bayesian Inference in Statistical Analysis*, (Wiley, New York, 1992).
2. M.-H. Chen, Importance-weighted marginal Bayesian posterior density estimation, *J. Amer. Statist. Assoc.* **89**, 818-824 (1994).
3. M.-H. Chen and Q.-M. Shao, Monte Carlo estimation of Bayesian credible and HPD intervals, *J. Comp. Graph. Statist.* **8**, 69-92 (1999).
4. M.-H. Chen, Q.-M. Shao, and J.G. Ibrahim, *Monte Carlo Methods in Bayesian Computation*, (Springer-Verlag, New York, 2000).
5. A.E. Gelfand, A.F.M. Smith, and T.M. Lee, (1992). Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling, *J. Amer. Statist. Assoc.* **87**, 523-532 (1992).
6. R.J. Hyndman, Computing and graphing highest density regions, *Amer. Statist.* **50**, 120-126 (1996).
7. B.W. Silverman, *Density Estimation for Statistics and Data Analysis*, (Chapman & Hall, London, 1986).
8. M.A. Tanner, *Tools for Statistical Inference.* Third Edition, (Springer-Verlag, New York, 1996).
9. G.C.G. Wei and M.A. Tanner, Calculating the content and boundary of the highest posterior density region via data augmentation. *Biometrika* **77**, 649-652 (1990).