

## Adaptive Gaussian Process Approximation for Bayesian Inference with Expensive Likelihood Functions

**Hongqiao Wang**

*wanghongqiao@sjtu.edu.cn*

*Institute of Natural Sciences and School of Mathematical Sciences,  
Shanghai Jiao Tong University, Shanghai 200240, China*

**Jinglai Li**

*jinglai.li@liverpool.ac.uk*

*Department of Mathematical Sciences, University of Liverpool, Liverpool, UK*

We consider Bayesian inference problems with computationally intensive likelihood functions. We propose a Gaussian process (GP)-based method to approximate the joint distribution of the unknown parameters and the data, built on recent work (Kandasamy, Schneider, & Póczos, 2015). In particular, we write the joint density approximately as a product of an approximate posterior density and an exponentiated GP surrogate. We then provide an adaptive algorithm to construct such an approximation, where an active learning method is used to choose the design points. With numerical examples, we illustrate that the proposed method has competitive performance against existing approaches for Bayesian computation.

### 1 Introduction ---

Bayesian inference is a popular method for estimating unknown parameters from data, and a major advantage of the method is its ability to quantify uncertainty in the inference results (Gelman et al., 2013). In this letter, we consider Bayesian inference problems where the likelihood functions are highly expensive to evaluate. A typical example of this type of problems is the Bayesian inverse problem (Tarantola, 2005), where the parameters of interest cannot be observed directly and need to be estimated from indirect data. Such problems arise from many real-world applications, ranging from carbon capture (Konomi, Karagiannis, Lai, & Lin, 2018) to chemical kinetics (Golightly & Wilkinson, 2011). In Bayesian inverse problems, the mappings from the parameter of interest to the observable quantities, often

---

J.L. is now at the University of Liverpool.

known as the forward models, are often computationally intensive (e.g., involving simulating large-scale computer models).

Due to the high computational cost, common numerical implementations of Bayesian inferences, such as Markov chain Monte Carlo (MCMC; Andrieu, De Freitas, Doucet, & Jordan, 2003) methods, can be prohibitively expensive. A simple idea to accelerate the computation of the posterior is to construct a computationally inexpensive surrogate or an approximation of the posterior distribution with a limited number of likelihood function evaluations. To this end, a convenient choice for surrogate function is the gaussian process (GP) model (Williams & Rasmussen, 2006). The idea of using this model to approximate the posterior or the likelihood function dates back to the so-called Bayesian quadrature (or Bayesian Monte Carlo) approaches (O'Hagan, 1991; Rasmussen & Ghahramani, 2003; Rasmussen et al., 2003; Kennedy, 1998), which were designed to perform numerical integrations in a Bayesian fashion (e.g., to compute the evidence in Bayesian inference problems; Osborne et al., 2012).

Unlike the Bayesian quadrature methods, the goal of this work is to construct an approximation of the posterior distribution. To this end, recent work (Kandasamy, Schneider, & Póczos, 2015) approximates the joint distribution of the unknown parameter and the data (which can also be viewed as the unnormalized posterior distribution) with an exponentiated GP model, where the design points (the points where the likelihood function is evaluated) are chosen with an active learning strategy. In particular, they determine the design points by sequentially maximizing the variance in the posterior approximation. Other ideas for using GP approximation to accelerate Bayesian computation can be found in Conrad, Marzouk, Pillai, and Smith (2016); Billionis and Zabaras (2013), among others. The method presented in this letter also intends to approximate the unnormalized posterior distribution. The main contribution of the work is as follows. We write the unnormalized posterior distribution as a product of an approximate posterior density and an exponentiated GP surrogate. The intuition behind this formulation is that the GP model can be more effectively constructed if we factor out a good approximation of the posterior (see section 2.3 for a detailed explanation). As we may not know a good approximate posterior density in advance, we develop an algorithm to adaptively construct the product-form approximation of the unnormalized posterior distribution. Another difference between our method and that in Kandasamy et al. (2015) is the learning strategy for selecting the design points. We use the entropy rather than the variance as the selection criterion, which can better represent the uncertainty in the approximation. Numerical examples illustrate that the proposed method can substantially improve the performance of the GP approximation.

We note that other surrogate models, notably generalized polynomial chaos (gPC) expansion (Li & Marzouk, 2014; Marzouk & Xiu, 2009; Marzouk & Najm, 2009; Marzouk, Najm, & Rahn, 2007; Nagel & Sudret,

2016), have also been used to accelerate Bayesian computation. Detailed comparison of the two types of the surrogates is not discussed in this work, and those who are interested in this matter may consult O'Hagan (2013).

The rest of the letter is organized as the following. In section 2, we present the adaptive GP algorithm to construct the posterior approximation and active learning method to determine the design points. In section 3, we give two examples to illustrate the performance of the proposed method. Section 4 provides some concluding remarks.

## 2 The Adaptive GP Method

---

**2.1 Problem Setup.** A Bayesian inference problem aims to estimate an unknown parameter  $\mathbf{x}$  from data  $\mathbf{d}$ . Specifically, it computes the posterior distribution of  $\mathbf{x}$  using the Bayes' formula,

$$\pi(\mathbf{x}|\mathbf{d}) \propto \pi(\mathbf{x}, \mathbf{d}) = l(\mathbf{d}|\mathbf{x})\pi(\mathbf{x}), \quad (2.1)$$

where  $l(\mathbf{d}|\mathbf{x})$  is the likelihood function and  $\pi(\mathbf{x})$  is the prior distribution of  $\mathbf{x}$ . When the Bayesian method is applied to inverse problems, the data and the forward model enter the formulation through the likelihood function. Suppose that there is a function (termed the *forward function* or *forward model*) that maps the parameter of interest  $\mathbf{x}$  to the observable quantity  $\mathbf{y}$ ,

$$\mathbf{y} = \mathbf{G}(\mathbf{x}) + \mathbf{z},$$

where  $\mathbf{z}$  is the observation error. Now we further assume that the distribution density of the observation noise  $\mathbf{z}$ ,  $p_z(\mathbf{z})$ , is available, and it follows directly that the likelihood function is given by

$$l(\mathbf{d}|\mathbf{x}) = p_z(\mathbf{d} - \mathbf{G}(\mathbf{x})).$$

In what follows, we omit the argument  $\mathbf{d}$  in the likelihood function and denote it as  $l(\mathbf{x})$  for simplicity. It is easy to see that each evaluation of the likelihood function  $l(\mathbf{x})$  requires evaluating the forward function  $\mathbf{G}(\mathbf{x})$ . In practice, the forward function  $\mathbf{G}(\mathbf{x})$  often represents a large-scale computer model, and thus the evaluation of  $l(\mathbf{x})$  can be computationally demanding. Due to the high computational cost, the brute force Monte Carlo simulation cannot be used for such problems, and we resort to an alternative method to compute the posterior distributions, using the GP surrogate model. A brief description of the GP method is provided in next section.

**2.2 The GP Model.** Given a real-valued function  $g(\mathbf{x})$ , the GP or the Kriging method constructs a surrogate model of  $g(\mathbf{x})$  in a nonparametric

Bayesian regression framework (Williams & Rasmussen, 2006; Oakley & O'Hagan, 2002; O'Hagan & Kingman, 1978). Specifically the target function  $g(\mathbf{x})$  is cast as a gaussian random process whose mean is  $\mu(\mathbf{x})$  and covariance is specified by a kernel function  $k(\mathbf{x}, \mathbf{x}')$ :

$$\text{COV}[g(\mathbf{x}), g(\mathbf{x}')] = k(\mathbf{x}, \mathbf{x}').$$

The kernel  $k(\mathbf{x}, \mathbf{x}')$  is positive semidefinite and bounded. Let us assume that  $m$  evaluations of the function  $g(\mathbf{x})$  are performed at parameter values  $\mathbf{X}^* := [\mathbf{x}_1^*, \dots, \mathbf{x}_m^*]$ , yielding function evaluations  $\mathbf{y}^* := [y_1^*, \dots, y_m^*]$ , where

$$y_i^* = g(\mathbf{x}_i^*) \quad \text{for } i = 1, \dots, m.$$

Suppose that we want to predict the function values at points  $\mathbf{D} := [\mathbf{x}_1, \dots, \mathbf{x}_{m'}]$ , that is,  $\mathbf{y} = [y_1, \dots, y_{m'}]$  where  $y_i = g(\mathbf{x}_i)$ . The sets  $\mathbf{X}^*$  and  $\mathbf{D}$  are often known as the training and the test points, respectively. The joint prior distribution of  $(\mathbf{y}^*, \mathbf{y})$  is

$$\begin{bmatrix} \mathbf{y}^* \\ \mathbf{y} \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mu(\mathbf{X}^*) \\ \mu(\mathbf{D}) \end{bmatrix}, \begin{bmatrix} K(\mathbf{X}^*, \mathbf{X}^*) & K(\mathbf{X}^*, \mathbf{D}) \\ K(\mathbf{D}, \mathbf{X}^*) & K(\mathbf{D}, \mathbf{D}) \end{bmatrix} \right), \quad (2.2)$$

where we use the notation  $K(\mathbf{A}, \mathbf{B})$  to denote the matrix of the covariance evaluated at all pairs of points in sets  $\mathbf{A}$  and  $\mathbf{B}$ . The posterior distribution of  $\mathbf{y}$  is also gaussian:

$$\mathbf{y} | \mathbf{D}, \mathbf{X}^*, \mathbf{y}^* \sim \mathcal{N}(\mathbf{u}, \Sigma), \quad (2.3a)$$

where the posterior mean is

$$\mathbf{u} = \mu(\mathbf{D}) + K(\mathbf{D}, \mathbf{X}^*)K(\mathbf{X}^*, \mathbf{X}^*)^{-1}(\mathbf{y}^* - \mu(\mathbf{X}^*)), \quad (2.3b)$$

and the posterior covariance matrix is

$$\Sigma = K(\mathbf{D}, \mathbf{D}) - K(\mathbf{D}, \mathbf{X}^*)K(\mathbf{X}^*, \mathbf{X}^*)^{-1}K(\mathbf{X}^*, \mathbf{D}). \quad (2.3c)$$

Here we provide only a brief introduction to the GP method tailored to our own purposes. Readers interested in further details may consult the references already mentioned.

**2.3 The Adaptive GP Algorithm.** Now we discuss how to use the GP method to compute the posterior distribution in our problem. A straightforward idea is to construct the surrogate model directly for the log-likelihood function  $\log l(\mathbf{x})$ ; such a method has been used in Osborne et al. (2012) and

Kandasamy et al. (2015). A difficulty in this approach is that the target function  $\log l(\mathbf{x})$  can be highly nonlinear and fast varying, and thus is not well described by a GP model. We present an adaptive scheme to alleviate the difficulty.

We first write the unnormalized posterior (i.e., the joint distribution  $\pi(\mathbf{x}, \mathbf{d})$ ) as

$$f(\mathbf{x}) = l(\mathbf{x})\pi(\mathbf{x}) = \exp(g(\mathbf{x}))p(\mathbf{x}),$$

where  $p(\mathbf{x})$  is a probability distribution that we are free to choose and

$$g(\mathbf{x}) = \log(f(\mathbf{x})/p(\mathbf{x})). \quad (2.4)$$

It is clear that equation 2.4 may become problematic when  $p(\mathbf{x}) = 0$  and  $f(\mathbf{x}) \neq 0$ , and to avoid the issue, we require that the support of  $p(\mathbf{x})$  be a superset of that of  $f(\mathbf{x})$ . We work on the log posterior distribution since the log smoothes out a function and is more conducive for the GP modeling. Also, by doing this, we ensure the nonnegativity of the obtained approximate posterior. We then sample the function  $g(\mathbf{x})$  at certain locations and construct the GP surrogate of  $g(\mathbf{x})$ . It should be noted that the distribution  $p(\mathbf{x})$  plays an important role in the surrogate construction, as a good choice of  $p(\mathbf{x})$  can significantly improve the accuracy of the GP surrogate models. In particular, if we take  $p(\mathbf{x})$  to be exactly the posterior  $\pi(\mathbf{x}|\mathbf{d})$ , it follows immediately that  $g(\mathbf{x})$  in equation 2.4 is a constant. This then gives us the intuition that if  $p(\mathbf{x})$  is a good approximation to the posterior distribution  $\pi(\mathbf{x}|\mathbf{d})$ ,  $g(\mathbf{x})$  is a mildly varying function that is easy to approximate. In other words, we can improve the performance of the GP surrogate by factoring out a good approximation of the posterior. Certainly this cannot be done in one step, as the posterior is not known in advance.

We present an adaptive framework to construct a sequence of pairs  $\{p_i(\mathbf{x}), \exp(\hat{g}_i(\mathbf{x}))\}$ , the product of which evolves to a good approximation of the unnormalized posterior  $f(\mathbf{x})$ . Roughly speaking the algorithm performs the following iterations: in the  $n$ th cycle, given the current guess of the posterior distribution  $p_n(\mathbf{x})$ , we construct a GP surrogate  $\hat{g}_n(\mathbf{x})$  of  $g_n(\mathbf{x})$ , which is given by

$$g_n(\mathbf{x}) = \log(f(\mathbf{x})/p_n(\mathbf{x})),$$

and we then compute a new (and possibly better) posterior approximation  $p_{n+1}(\mathbf{x})$  using

$$p_{n+1}(\mathbf{x}) \propto \exp(\hat{g}_n(\mathbf{x}))p_n(\mathbf{x}).$$

Finally we want to specify stopping criteria for the iteration, and the iteration terminates if either of the following two conditions is satisfied. The first is that the maximum number of iterations is reached. Our second stopping condition is based on the Kullback-Leiber divergence (KLD) between  $p_{n-1}$  and  $p_n$ , which reads,

$$D_{KL}(p_{n-1}, p_n) = \int \log \frac{p_{n-1}(\mathbf{x})}{p_n(\mathbf{x})} p_{n-1}(\mathbf{x}) d\mathbf{x}. \quad (2.5)$$

Specifically the second stopping condition is that  $D_{KL}(p_{n-1}, p_n)$  is smaller than a prescribed value  $D_{\max}$  in  $K$  consecutive iterations. That is, if the computed posterior approximation does not change much in a certain number of consecutive iterations, the algorithm terminates. The complete scheme is described in algorithm 1.

Some remarks on the implementation of algorithm 1 are listed in order:

- In line 6, we construct the GP model for  $g_n(\mathbf{x})$  using the procedure described in section 2.2. The hyperparameters of the GP model are determined by maximizing the marginal likelihood function (Williams & Rasmussen, 2006).
- In line 7, we resort to the MCMC method to draw a rather large number of samples from the approximate posterior distribution. This procedure, however, does not require evaluating the true likelihood function and is not computationally expensive.
- In line 8, we need to compute the density function of a distribution  $p_{n+1}$  from the samples  $X_n$ , and here we use the gaussian mixture method (McLachlan & Peel, 2004) to estimate the density. We note that an important issue in the gaussian mixture method is how to choose the number of mixtures. In our numerical examples, we use a fixed number of mixtures, but it is also possible to determine it with an adaptive procedure (Vlassis & Likas, 2002). Certainly there will be estimation errors in this procedure and so we denote the estimated density as  $\hat{p}_{n+1}$  to distinguish it from the true density  $p_{n+1}$ .
- In line 9, we find that it is rather costly to compute the KLD between  $p_{n-1}$  and  $p_n$ . We instead use the KLD between  $\hat{p}_{n-1}$  and  $\hat{p}_n$ , which is much easier to compute as the distributions are available as gaussian mixtures. The KLD can be efficiently computed with the Monte Carlo method as we have stored the samples  $A_{n-1}$  from  $\hat{p}_{n-1}$  in the previous iteration.
- In line 15, we need to determine the design points—the locations where we evaluate the true function. The choice of design points is critical to the performance of the proposed adaptive GP algorithm, and we use an active learning method to determine the points, which we present in section 2.4.

---

**Algorithm 1:** The Adaptive GP Algorithm.
 

---

**Require:**  $m_0, n_{\max}, M, D_{\max}, K$ 

- 1: let  $\hat{p}_0(x) = \pi(x)$ ; let  $n = 0$ ; let  $k = 0$ ;
- 2: choose  $m_0$  initial design points:  $\{x_1, \dots, x_{m_0}\}$ , and compute  $y_i = f(x_i)$  for  $i = 1 \dots m_0$ ;
- 3: let  $S_0 = \{(x_1, y_1), \dots, (x_{m_0}, y_{m_0})\}$ ;
- 4: **for**  $n=0$  **do**  $n_{\max}$
- 5:     let  $g_n(x) = \log(f(x)/\hat{p}_n(x))$ ;
- 6:     construct a GP surrogate model  $\hat{g}_n(x)$  for the function  $g_n(x)$  with data set  $S_n$ ;
- 7:     draw a set of  $M$  samples from the approximate posterior

$$p_{n+1}(x) \propto \exp(\hat{g}_n(x))\hat{p}_n(x)$$

 with MCMC, denoted as  $A_n$ ;

- 8:     obtain an estimated PDF from samples  $A_n$ , denoted as  $\hat{p}_{n+1}$ ;
  - 9:     compute  $D_{KL}(\hat{p}_{n-1}, \hat{p}_n)$ ;
  - 10:    **if**  $D_{KL}(\hat{p}_{n-1}, \hat{p}_n) < D_{\max}$  **then**  $k = k + 1$ ;
  - 11:    **else**  $k = 0$ ;
  - 12:    **end if**
  - 13:    **if**  $k = K$  **then** break the FOR loop;
  - 14:    **else**
  - 15:       select  $m$  design points:  $\{x_1, \dots, x_m\}$ , evaluate  $f(x_i)$  for  $i = 1 \dots m$ , and let
 
$$S_{n+1} = S_n \cup \{(x_1, y_1), \dots, (x_m, y_m)\}$$
;
  - 16:    **end if**
  - 17: **end for**
-

**2.4 Active Learning for the Design Points.** In the GP literature, the determination of the design points is often cast as an experimental design problem, that is, we need to find the experimental parameters that can provide us the most information. The problem has received considerable attention, and a number of methods and criteria have been proposed to select the points: the mutual information criterion (Krause, Singh, & Guestrin, 2008), the integrated mean square error (IMSE; Sacks, Welch, Mitchell, & Wynn, 1989), the integrated posterior variance (IVAR; Gorodetsky & Marzouk, 2016), and the active learning MacKay (ALM; criterion MacKay, 1992), to name just a few. Here we choose to use an active learning strategy that adds one design point a time, primarily because it is easy to implement.

A common active learning strategy is to choose the point that has the largest uncertainty, and to this end we need a function that can measure or quantify the uncertainty in the approximation reconstructed. In the usual GP problems, the variance of the GP model  $\hat{g}(\mathbf{x})$  is a natural choice for such a measure of uncertainty (which yields the ALM method), because the distribution of  $\hat{g}(\mathbf{x})$  is gaussian. In our problems, however, the function of interest is the posterior approximation  $\hat{f}(\mathbf{x}) = \exp(\hat{g}(\mathbf{x}))p(\mathbf{x})$  rather than the GP model  $\hat{g}(\mathbf{x})$  itself, and thus we should measure the uncertainty in  $\hat{f}(\mathbf{x})$ . In Kandasamy et al. (2015), the variance of the posterior approximation  $\hat{f}$  is used as the measure function. However, since the distribution of  $\hat{f}(\mathbf{x})$  is not gaussian, the variance may not provide a good estimate of the uncertainty. On the other hand, the entropy is a commonly used measure to quantify the uncertainty in a random variable (Rényi, 1961; Shannon, 2001), and here we use it as our design criterion.

Specifically, suppose that at point  $\mathbf{x}$ , the distribution of  $\hat{f}(\mathbf{x})$  is  $\pi_f(\hat{f})$ , and the entropy of  $\hat{f}(\mathbf{x})$  is defined as

$$\mathbb{H}(\hat{f}(\mathbf{x})) = - \int \log(\pi_f(\hat{f}))\pi_f(\hat{f})d\hat{f}. \quad (2.6)$$

Thus, we choose a new design point by

$$\max_{\mathbf{x} \in \Omega} \mathbb{H}(\hat{f}(\mathbf{x})),$$

where  $\Omega$  is a bounded subspace of the state space of  $\mathbf{x}$ . In the present problem, the distribution of  $\hat{g}(\mathbf{x})$  is gaussian, and let us assume its mean and variance are  $\mu$  and  $\sigma^2$ , respectively. It follows that the distribution of  $\hat{f}(\mathbf{x})$  is log normal and the entropy of it can be computed analytically:

$$\mathbb{H}(\hat{f}) = \mu + \frac{1}{2} \ln(2\pi e\sigma^2) + \log p. \quad (2.7)$$



We want to emphasize that the entropy-based active learning method is different from the usual maximum entropy method for experimental design, (Sebastiani & Wynn, 2000). The purpose of the maximum entropy method in Sebastiani and Wynn (2000) is to find the design points that maximize the information gain of an inference problem, while in our problem, we use entropy as a measure of uncertainty.

Now suppose that we have a set of existing data points, and we want to choose  $m$  new design points. We use the following scheme to sequentially choose the new points:

1. Construct a GP model  $\hat{g}(\mathbf{x})$  for  $g(\mathbf{x})$  using data set  $S$ .
2. Compute  $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \Omega} \mathbb{H}(\hat{f}(\mathbf{x}))$ .
3. Evaluate  $y^* = g(\mathbf{x}^*)$  and let  $S = S \cup \{(\mathbf{x}^*, y^*)\}$ .

Note that the key in the adaptive scheme is step 2, where we seek the point  $\mathbf{x}$  that maximizes the entropy  $\mathbb{H}(\hat{f}(\mathbf{x}))$  in  $\Omega$ . This is a challenging problem from an optimization perspective, because the problem may have multiple local maxima. However, in the numerical tests, we have found that our algorithm does not strictly require the optimality of the solution, and it performs well as long as a good design point can be found in each step. Thus, here we use a stochastic search method, the simulated annealing algorithm (Kirkpatrick, Gelatt, & Vecchi, 1983), to find the design point. We have tested other metaheuristic optimization algorithms, and the performances do not vary significantly.

### 3 Numerical Examples

---

**3.1 The Rosenbrock Function.** We first test our method on a two-dimensional mathematical example. The likelihood function is

$$l(\mathbf{x}) = \exp \left( -\frac{1}{100}(x_1 - 1)^2 - (x_1^2 - x_2)^2 \right), \quad (3.1)$$

which is the well-known Rosenbrock function, and the prior  $\pi(\mathbf{x})$  is a uniform distribution defined on  $[-5, 5] \times [-5, 5]$ . The resulting unnormalized posterior is shown in Figure 1 (left). The function has a banana shape and is often used as a test problem for Bayesian computation methods.

We now apply the proposed adaptive GP method to compute the posterior for this problem. In this example, we let  $m_0 = 20$ , and the samples in  $S_0$  were randomly drawn according to the prior distribution. We also choose  $m = 10$ : 10 new design points are computed in each iteration. In the algorithm, we need to sample from the approximate posterior distribution in each iteration, and here we draw  $M = 2 \times 10^4$  samples with the delayed rejection adaptive Metropolis algorithm (DRAM; Haario, Laine, Mira, & Saksman, 2006). We restate that the  $2 \times 10^4$  MCMC samples are generated from




Figure 1: (Left) True posterior distribution. (Right) KL distance between  $p_{n-1}$  and  $p_n$ , plotted against the number of iterations.

the approximate posterior distribution, and thus it does not require evaluating the true likelihood function. We also set the parameters that specify the termination conditions to be  $n_{\max} = 100$ ,  $D_{\max} = 0.01$ , and  $K = 5$ . The algorithm terminates with 13 iterations, and a total of 140 evaluations of the true likelihood function are used. In Figure 1 (right), we plot the KL difference in two consecutive iterations, which is used as one of our stopping criteria, against the number of iterations. To illustrate the performance of our method, we use the KL distance and the Hellinger distance, defined as

$$D_H(p_1, p_2) = \frac{1}{2} \int (\sqrt{p_1(\mathbf{x})} - \sqrt{p_2(\mathbf{x})})^2 dx,$$

to quantify the difference between the computed approximation and the true posterior. We plot the KL (left) and the Hellinger (right) distances between the approximate posterior and the true posterior distribution in Figure 2. It can be seen from the figures that the computed approximation converges very well to the true posterior in terms of both distance measures as the iteration proceeds. We then plot the approximate posterior obtained in the 7th, 9th, 11th, and 13th iterations in Figure 3, in which we can visualize how the quality of the approximation increases as the iterations proceed. In each of the plots, we also show the design points (red dots) that have been used up to the given iteration. As a side product, we can obtain the variance of the approximate posterior, which may provide the uncertainty information of the approximation, and here we plot the variance in Figure 4. As a comparison, we also compute the GP approximation of the posterior with the Bayesian active posterior estimation (BAPE) method developed in Kandasamy et al. (2015). In particular, we implement the BAPE method using a total of 140 design points, which matches the number of design points of our method. The results are shown in Figure 5. The figure on the left shows the




Figure 2: KL (left) and the Hellinger (right) distances between the obtained approximation and the true posterior, plotted against the number of iterations.

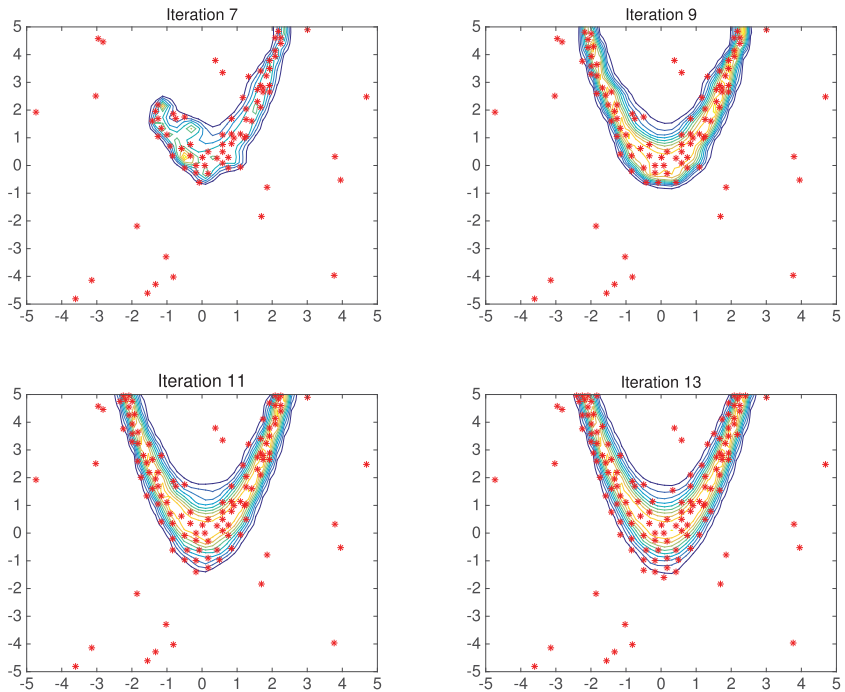


Figure 3: Approximate posterior distribution obtained at the 7th, 9th, 11th, and 13th iterations respectively. The red dots are the design points that have been used.




Figure 4: Variance of the posterior distribution obtained with the AGP method.




Figure 5: (Left) GP approximation of the posterior distribution obtained with the BAPE method using 140 design points (red dots). (Right) KL distance between the true posterior and the approximation computed with the AGP (solid line) and the BAPE (dashed line) methods, plotted against the number of design points used. The inset is the same plot on a logarithmic scale.

posterior distribution computed with all 140 design points (corresponding to the 13th iteration in our method), and as one can see, the BAPE method can also obtain a good approximation of the posterior distribution. To compare the performance of the two methods, we compute the KLD between the true posterior and the approximation obtained with different numbers of design points by the BAPE and our adaptive GP (AGP) methods. We plot

the KL distance against the number of design points in Figure 5 (left). One can see from the figure that with the same number of design points, the approximate posterior obtained by the proposed AGP method is significantly closer to the true posterior than the results of the BAPE method.

**3.2 Genetic Toggle Switch.** We now apply the proposed method to a real-world inference problem; we consider the kinetics of a genetic toggle switch, which was first studied in Gardner, Cantor, and Collins (2000) and later numerically investigated in Marzouk and Xiu (2009). The toggle switch consists of two repressible promoters arranged in a mutually inhibitory network: promoter 1 and promoter 2. Either promoter transcribes a repressor for the other one; moreover, either repressor may be induced by an external chemical or thermal signal. Genetic circuits of this form can be modeled by the following differential-algebraic equation system (Gardner et al., 2000):

$$\frac{du}{dt} = \frac{\alpha_1}{1 + v^\beta} - u, \quad (3.2a)$$

$$\frac{dv}{dt} = \frac{\alpha_2}{1 + w^\gamma} - v, \quad (3.2b)$$

$$w = \frac{u}{1 + ([IPTG]/K)^\eta}. \quad (3.2c)$$

In equations 3.2a to 3.2c  $u$  and  $v$  are, respectively, the concentration of repressors 1 and 2;  $\alpha_1$  and  $\alpha_2$  are the effective rates of synthesis of the repressors;  $\gamma$  and  $\beta$  represent cooperativity of repression of the two promoters; and [IPTG] is the concentration of IPTG, the chemical compound that induces the switch. Parameters  $K$  and  $\eta$  describe binding of IPTG with the first repressor. (For more details of the model, see Gardner et al., 2000.)

We performed the experiments with several selected values of [IPTG]:  $1 \times 10^{-6}$ ,  $5 \times 10^{-4}$ ,  $7 \times 10^{-4}$ ,  $1 \times 10^{-3}$ ,  $3 \times 10^{-3}$ ,  $5 \times 10^{-3}$ , respectively, and for each experiment, we look the measurement of  $v$  at  $t = 10$ . The goal is to infer the six parameters,

$$\mathbf{x} = [\alpha_1, \alpha_2, \gamma, \beta, \eta, K],$$

from the measurements of  $v$ . We use synthetic data in this problem; specifically we assume that the true values of the parameters are

$$\mathbf{x}_{\text{true}} = [143, 15.95, 2.70, 0.96, 2.34, 2.70 \times 10^{-5}].$$

The data are simulated using the model described by equations 3.2a to 3.2c, with the true parameter values and measurement noise then added to the simulate data. The measurement noise here is assumed gaussian and zero-mean, with a variance  $\sigma^2$ . In the numerical experiments, we

Table 1: Prior Domains of the Parameters.

	$\alpha_1$	$\alpha_2$	$\gamma$	$\beta$	$\eta$	$K$
$n$	[120, 200]	[15.0, 16.0]	[2.1, 2.9]	[0.85, 1.15]	[1.3, 2.7]	$[2.3, 3.7] \times 10^{-5}$

consider a large-noise case where  $\sigma^2 = 5 \times 10^{-4}$  and a small-noise case where  $\sigma^2 = 1.25 \times 10^{-4}$ . We assume that the priors of the six parameters are all uniform and independent of each other, where the domains of the uniform priors are given in Table 1.

We use this example to make a detailed comparison of the proposed AGP algorithm with some other popular methods. Thus, we employ four methods to compute the posterior distribution in this example: the direct MCMC algorithm with the true likelihood function, the proposed AGP algorithm, the BAPE method (Kandasamy et al., 2015), and the spectral likelihood expansion (SLE) method (Nagel & Sudret, 2016), which constructs the gPC surrogate for the likelihood function using nonintrusive approaches.

We first consider the large-noise case. We draw  $3 \times 10^5$  samples from the true posterior distribution with a DRAM algorithm and use the results as the reference posterior distribution. We then apply the AGP method to approximate the posterior distribution, where we use  $m_0 = 50$  initial design points randomly drawn from the prior and  $m = 50$  design points in each iteration. We also choose the termination parameters to be  $n_{\max} = 100$ ,  $D_{\max} = 0.05$ , and  $K = 5$ . The algorithm terminates in 18 iterations, resulting in a total of 950 evaluations of the true likelihood function. We note that each evaluation of the likelihood function involves a full simulation of the underlying model described by equations 3.2a to 3.2c. After obtaining the approximate posterior distribution, we draw  $3 \times 10^5$  samples from it using a DRAM MCMC simulation.

We then compute the approximate posterior with the BAPE method using 950 likelihood evaluations and draw  $3 \times 10^5$  samples from it using a DRAM MCMC simulation. Finally, we approximate the posterior with the SLE-gPC method where the gPC expansion coefficients are computed using the least squares method with design points determined by the Sobol sequences–based quasi–Monte Carlo (QMC) method. The gPC degree is automatically determined by the algorithm using leave-one-out (LOO) cross-validation, and in the QMC scheme, we set the number of design points at 950.

We now compare the results of these methods. First, we estimate the posterior distributions of the six parameters by all four methods and show the results in Figure 6. One can see from the figure that the distributions computed by both the BAPE and the AGP methods are rather close to those of direct MCMC (which are regarded as the true posteriors), while the results of SLE-gPC deviate from the MCMC results, especially for the two

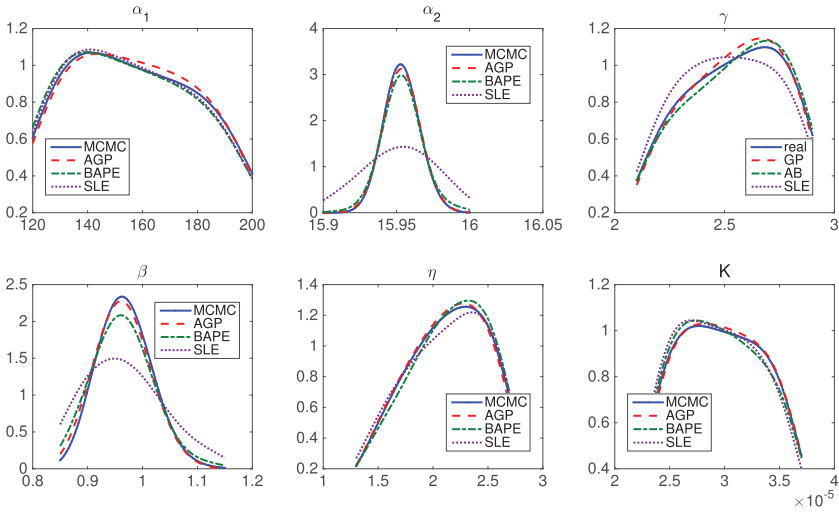


Figure 6: Genetic toggle example: large-noise case. The marginal distributions of the six parameters computed with the four different methods.

Table 2: Large-Noise Case: The KLD between the Marginal Posterior Distributions Computed with the Three Approximate Methods and Those Computed with Standard MCMC.

	Method	$\alpha_1$	$\alpha_2$	$\gamma$	$\beta$	$\eta$	$K$
KLD	AGP	$1.5 \times 10^{-4}$	0.0069	0.0020	0.014	0.0041	0.0075
	BAPE	$4.5 \times 10^{-3}$	0.043	0.0022	0.039	0.0095	0.013
	SLE	$1.8 \times 10^{-3}$	0.44	0.0012	0.18	0.0078	0.0013

parameters  $\alpha_2$  and  $\beta$ . As for the comparison of BAPE and AGP, the figures show that both methods can produce reasonably good approximations of the posterior distributions in this case. So for a quantitative evaluation of the performance of the methods, we compute the KLD from the approximate posterior distributions to the true posterior densities and show the results in Table 2. We can see from the table that the posterior distributions computed by the AGP method are closer (in terms of KLD) to the true posteriors than the other two methods for all six parameters.

We then consider the small-noise case where we use the same implementation configurations as in the large-noise case. In this case, the AGP algorithm uses 1200 true likelihood evaluations; as before, we compute the posterior using the SLE and the BAPE methods with the same number of true likelihood evaluations. We then compute the posterior directly with  $3 \times 10^5$  MCMC samples and use the results as the true posterior. In Figure 7

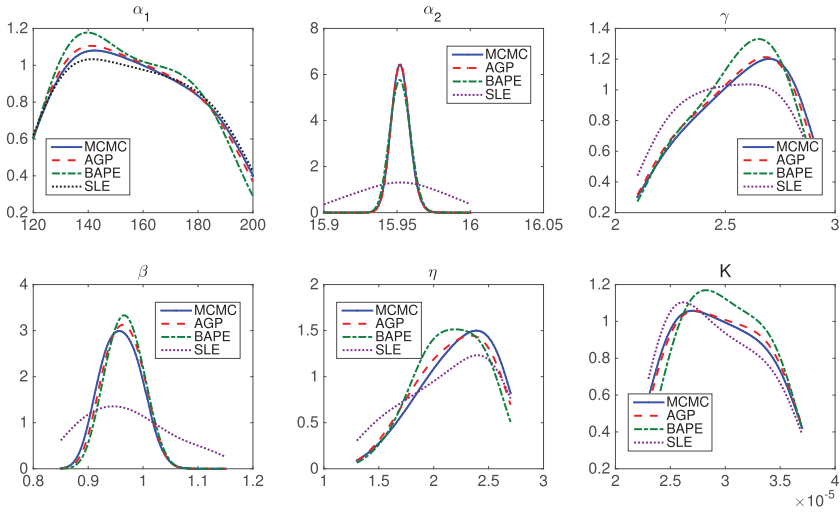


Figure 7: Genetic toggle example for the small-noise case. The marginal distributions of the six parameters, computed with the four different methods.

Table 3: Small-Noise Case: The KLD between the Marginal Posterior Distributions Computed with the Three Approximate Methods and Those Computed with Standard MCMC.

	Method	$\alpha_1$	$\alpha_2$	$\gamma$	$\beta$	$\eta$	$K$
KLD	AGP	0.0032	0.015	0.0039	0.012	0.0096	0.014
	BAPE	0.015	0.036	0.028	0.057	0.0075	0.074
	SLE	0.0035	1.1	0.021	0.46	0.050	0.010

we compare the marginal posterior distributions with the four methods. Similar to the large-noise case, the figure shows that the results of the SLE-gPC are of very low accuracy, while both the AGP and the BAPE methods yield rather good results. Once again, we show in Table 3 the KLD from the marginal posterior distributions computed with the three approximate methods to the true posterior (those computed by the direct MCMC). These quantitative comparison results indicate that the AGP method yields better results than the other two methods in terms of the KLD. Thus, we can conclude that our AGP method has the best performance in both the large- and the small-noise cases.

**3.3 The Human Body Sway Problem.** Finally, we apply the proposed method to a human body sway problem. This problem has received considerable attention as body sway may provide information about a



person’s physiological status (Tietäväinen, Gutmann, Keskivakkuri, Corander, & Häggström, 2017). Several mathematical models have been proposed to describe the sway motion, and here we consider the single-link inverted pendulum (SLIP) model proposed in Asai et al. (2009), which assumes that the body is maintained in an upright position by an active and a passive proportional-derivative controller.

Specifically, the SLIP model is given by the following stochastic delay differential equation (SDDE) (Asai et al., 2009):

$$I\ddot{\theta}(t) = mgh\theta(t) - [K\theta(t) + B\dot{\theta}(t) + f_P(\theta(t - \Delta)) + f_D(\dot{\theta}(t - \Delta))] + \xi(t). \tag{3.3}$$

In this equation,  $I$  is the moment of inertia of the body,  $\theta$  is the tilt angle ( $\dot{\theta}$  and  $\ddot{\theta}$  are its first and second derivatives, respectively),  $m$  denotes the body mass,  $g$  is the gradational acceleration,  $h$  is the distance between 3D center-of-mass (COM) and the ankle joint, and  $\xi$  is a zero-mean gaussian noise with variance  $\sigma^2$ .  $K$  and  $P$  are the passive stiffness and passive damping parameters, and  $f_P(\theta(t - \Delta))$  and  $f_D(\dot{\theta}(t - \Delta))$  are active stiffness and active damping terms, where  $\Delta$  is the time delay. We now specify the active stiffness  $f_P(\theta(t - \Delta))$  and the active damping  $f_D(\dot{\theta}(t - \Delta))$ . We first define two functions  $c_1(\theta(t - \Delta)) = \theta(\dot{\theta}(t - \Delta) - a_s\theta(t - \Delta))$  and  $c_2(\theta(t - \Delta)) = \theta(t - \Delta)^2 + (\dot{\theta}(t - \Delta))^2$ . We then have

$$f_P(\theta(t - \Delta)) = \begin{cases} P\theta(t - \Delta), & \text{if } c_1(\theta(t - \Delta)) > 0 \text{ and } c_2(\theta(t - \Delta)) > r^2; \\ 0, & \text{otherwise;} \end{cases} \tag{3.4}$$

and

$$f_D(\dot{\theta}(t - \Delta)) = \begin{cases} D\dot{\theta}(t - \Delta) & \text{if } c_1(\theta(t)) > 0 \text{ and } c_2(\theta(t)) > r^2; \\ 0, & \text{otherwise;} \end{cases} \tag{3.5}$$

where  $r$  is the radius of the quiet zone (active control is off). The slope  $a_s$  depends on the level of control,  $C_{ON}$ , as  $a_s = -\tan(\pi(C_{ON} - 0.5))$ . In this model, five key parameters ( $P$ ,  $D$ ,  $\Delta$ ,  $\sigma$ , and  $C_{ON}$ ) cannot be measured directly and need to be inferred from the body sway measurements, while the other parameters can either be measured or specified in advance (Tietäväinen et al., 2017). The COM signal,

$$\text{COM}(t) = h \sin(\theta(t)), \tag{3.6}$$

is measured and used to infer the five unknown model parameters. The absolute value of the COM amplitude, velocity, acceleration, and the power




Figure 8: Simulated COM signal.

spectral density (PSD) are extracted from the signal. The mean, variance, skewness, and kurtosis of each physical quantity are calculated as the data  $\mathbf{y}$  (a 16-dimensional vector) to infer the model parameters:  $\mathbf{x} = (P, D, \Delta, \sigma, C_{ON})$ . Computing the posterior in this problem is rather challenging as the likelihood function  $p(\mathbf{y}|\mathbf{x})$  is not available, which makes the standard Bayesian inference computation methods such as the MCMC algorithms infeasible. In Tietäväinen et al. (2017), the parameters were inferred with the approximate Bayesian computation (ABC) method, which does not use the likelihood function. A major issue in the ABC method is that its performance is often sensitive to the choice of the tolerance parameter (Barber, Voss, & Webster, 2015). Here as an alternative, we compute the posterior with the proposed method. In particular, we compute the likelihood function using the following procedure. For a given parameter value  $\mathbf{x}$ , we perform a Monte Carlo simulation for the SDDE model, equation 3.3, with a given number of samples. For each  $y_i$  for  $i = 1, \dots, 16$ , we estimate the resulting conditional density function  $p_i(y_i|\mathbf{x})$  with the kernel density estimation method, and then we take the likelihood function to be

$$p(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{16} p_i(y_i|\mathbf{x}).$$

We note that a single evaluation of the likelihood function requires repeatedly simulating equation 3.3 a large number of times, which renders the evaluation highly intensive.

In the numerical experiments, we use simulated data; in particular, the true parameter values are set to be  $P = 145$  Nm/rad,  $D = 10$  Nms/rad,  $\Delta = 0.2$  s,  $\sigma = 0.45$  Nm and  $C_{ON} = 0.75$ , and the other parameter values are  $g = 9.81$  m/s<sup>2</sup>,  $m = 68$  kg,  $h = 0.87$  m,  $I = mh^2$ ,  $K = mgh \times 0.8$  Nm/rad,  $B = 4$  Nms/rad, and  $r = 0.004$  rad-rad/s. The COM signal generated from the model, equation 3.3, with these parameter values is shown in Figure 8. In the inference, we impose a uniform distribution on each of the five parameters on the following intervals:  $P \in [80, 160]$ ,  $D \in [0.05, 30]$ ,  $\Delta \in [0.05, 0.5]$ ,  $\sigma \in [0.05, 0.6]$ , and  $C_{ON} \in [0.05, 0.85]$ . Moreover, for each evaluation of the

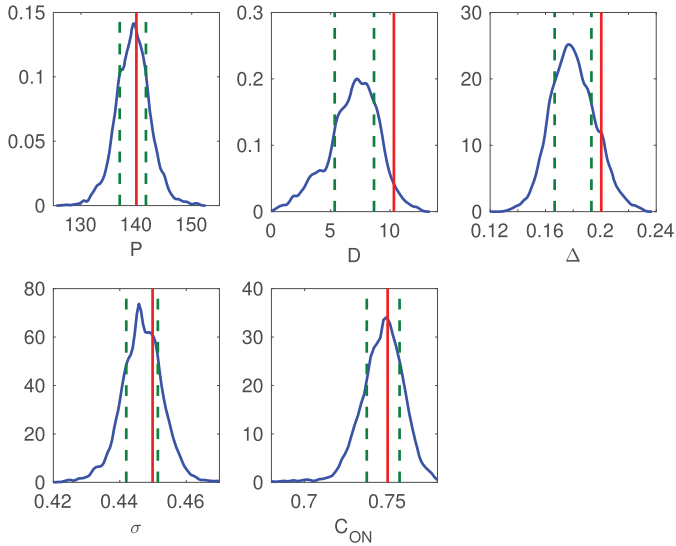


Figure 9: Posteriors of the parameters in the SLIP model, computed by the BAPE method. Also shown are the 60% confidence interval (dashed vertical lines) and the true parameter values (solid vertical lines).

likelihood function  $\pi(\mathbf{y}|\mathbf{x})$ , we use 10,000 simulations of equation 3.3, and as a result, a direct MCMC simulation of the posterior distribution is computationally infeasible. We apply our AGP method to compute the posterior distribution, and the algorithm parameters are the same as those in the second example. The algorithm terminates in 14 iterations, so the total number of true likelihood evaluations is 750. As a comparison, we also perform the BAPE method with the same number of true likelihood function evaluations. To compare the performance of the two methods, we plot the posterior marginals of the model parameters computed by BAPE in Figure 9 and those computed by AGP in Figure 10. We also show the true parameter values, as well as the 60% confidence interval in the figures. Here, one can see that for all the posteriors computed with the AGP method, the true parameter values fall in the 60% confidence intervals, while for the results of the BAPE method, the true values of  $D$  and  $\Delta$  fall outside the 60% confidence intervals, which suggests that the posteriors computed by the AGP method may be more accurate and reliable than those by BAPE.

#### 4 Conclusion

---

We have proposed an algorithm to construct GP-based approximations for the unnormalized posterior distribution. The method expresses the

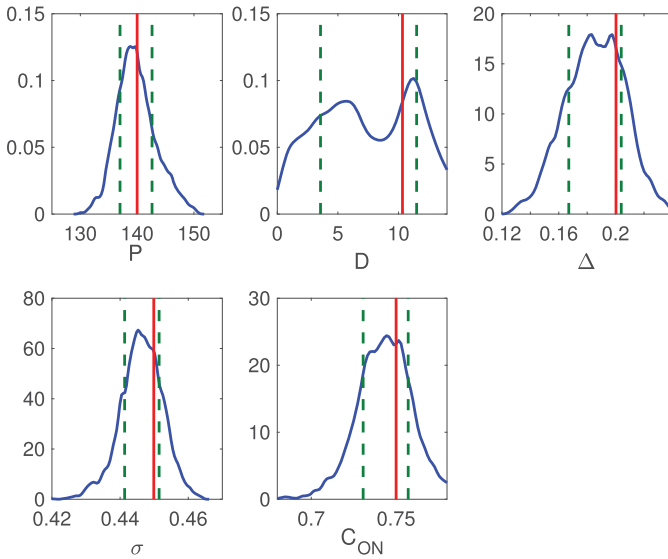


Figure 10: Posteriors of the parameters in the SLIP model, computed by our AGP method. Also shown are the 60% confidence interval (dashed vertical lines) and the true parameter values (solid vertical lines).

unnormalized posterior as a product of an approximate posterior density and an exponentiated GP model, and an adaptive scheme is presented to construct such an approximation. We also provide an active learning method that uses maximum entropy as the selection criterion to determine the sampling points. With numerical examples, we show that the method can obtain a rather good approximation of the posterior with a limited number of evaluations of the likelihood functions. We believe the proposed method can be useful in a wide range of practical Bayesian inference problems where the likelihood functions are difficult or expensive to evaluate.

Several issues of the proposed algorithm deserve further study. First, while our numerical experiments illustrate that the algorithm may converge in these examples, a rigorous convergence analysis of the algorithm is lacking. Second, for a posterior distribution with an unbounded domain, the resulting approximation may become improper, and thus certain modifications of the algorithm may be needed to address the issue. Finally, we note that selecting a good kernel function for the GP model is an important issue for GP-based methods, and to this end, an interesting question is how to choose kernel functions that are specifically suitable for the log posteriors. We plan to study these issues in future work.

## Acknowledgments

---

The work was partially supported by the National Natural Science Foundation of China under grant 11771289.

## References

---

- Andrieu, C., De Freitas, N., Doucet, A., & Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Machine Learning*, 50(1–2), 5–43.
- Asai, Y., Tasaka, Y., Nomura, K., Nomura, T., Casadio, M., & Morasso, P. (2009). A model of postural control in quiet standing: Robust compensation of delay-induced instability using intermittent activation of feedback control. *PloS One*, 4(7), e6169.
- Barber, S., Voss, J., & Webster, M. (2015). The rate of convergence for approximate Bayesian computation. *Electronic Journal of Statistics*, 9(1), 80–105.
- Bilionis, I., & Zabaras, N. (2013). Solution of inverse problems with limited forward solver evaluations: A Bayesian perspective. *Inverse Problems*, 30(1), 015004.
- Conrad, P. R., Marzouk, Y. M., Pillai, N. S., & Smith, A. (2016). Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *Journal of the American Statistical Association*, 111(516), 1591–1607.
- Gardner, T. S., Cantor, C. R., & Collins, J. J. (2000). Construction of a genetic toggle switch in *Escherichia coli*. *Nature*, 403(6767), 339–342.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). London: Chapman & Hall/CRC.
- Golightly, A., & Wilkinson, D. J. (2011). Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6), 807–820.
- Gorodetsky, A., & Marzouk, Y. (2016). Mercer kernels and integrated variance experimental design: Connections between gaussian process regression and polynomial approximation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1), 796–828.
- Haario, H., Laine, M., Mira, A., & Saksman, E. (2006). DRAM: Efficient adaptive MCMC. *Statistics and Computing*, 16(4), 339–354.
- Kandasamy, K., Schneider, J., & Póczos, B. (2015). Bayesian active learning for posterior estimation. In *Proceedings of the 24th International Conference on Artificial Intelligence* (pp. 3605–3611).
- Kennedy, M. (1998). Bayesian quadrature with non-normal approximating functions. *Statistics and Computing*, 8(4), 365–375.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671–680.
- Konomi, B. A., Karagiannis, G., Lai, K., & Lin, G. (2018). Bayesian treed calibration: An application to carbon capture with ax sorbent. *Journal of the American Statistical Association*, 112(517), 37–53.
- Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9, 235–284.

- Li, J., & Marzouk, Y. M. (2014). Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing*, 36(3), A1163–A1186.
- MacKay, D. J. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4(4), 590–604.
- Marzouk, Y. M., & Najm, H. N. (2009). Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics*, 228(6), 1862–1902.
- Marzouk, Y. M., Najm, H. N., & Rahn, L. A. (2007). Stochastic spectral methods for efficient Bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2), 560–586.
- Marzouk, Y., & Xiu, D. (2009). A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics*, 6(4), 826–847.
- McLachlan, G., & Peel, D. (2004). *Finite mixture models*. Hoboken, NJ: Wiley.
- Nagel, J. B., & Sudret, B. (2016). Spectral likelihood expansions for Bayesian inference. *Journal of Computational Physics*, 309, 267–294.
- Oakley, J., & O’Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4), 769–784.
- O’Hagan, A. (1991). Bayes–Hermite quadrature. *Journal of Statistical Planning and Inference*, 29(3), 245–260.
- O’Hagan, A. (2013). Polynomial chaos: A tutorial and critique from a statistician’s perspective. *SIAM/ASA Journal on Uncertainty Quantification Uncertainty Quantification*, 20, 1–20.
- O’Hagan, A., & Kingman, J. (1978). Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40, 1–42.
- Osborne, M., Garnett, R., Ghahramani, Z., Duvenaud, D. K., Roberts, S. J., & Rasmussen, C. E. (2012). Active learning of model evidence using Bayesian quadrature. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 25 (pp. 46–54). Red Hook, NY: Curran.
- Rasmussen, C. E., Bernardo, J., Bayarri, M., Berger, J., Dawid, A., Heckerman, D., . . . West, M. (2003). Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals. In J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, & M. West (Eds.), *Bayesian Statistics: 7*, 651–659. Oxford: Oxford University Press.
- Rasmussen, C. E., & Ghahramani, Z. (2003). Bayesian Monte Carlo. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, 16 (pp. 505–512). Cambridge, MA: MIT Press.
- Rényi, A. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1: *Contributions to the Theory of Statistics*. Berkeley: University of California Press.
- Sacks, J., Welch, W. J., Mitchell, T. J., & Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science*, 4, 409–423.
- Sebastiani, P., & Wynn, H. P. (2000). Maximum entropy sampling and optimal Bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1), 145–157.
- Shannon, C. E. (2001). A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1), 3–55.

- Tarantola, A. (2005). *Inverse problem theory and methods for model parameter estimation*. Philadelphia: SIAM.
- Tietäväinen, A., Gutmann, M. U., Keskivakkuri, E., Corander, J., & Hægström, E. (2017). Bayesian inference of physiologically meaningful parameters from body sway measurements. *Scientific Reports*, 7(1), 3771.
- Vlassis, N., & Likas, A. (2002). A greedy EM algorithm for gaussian mixture learning. *Neural Processing Letters*, 15(1), 77–87.
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.

---

Received March 19, 2017; accepted May 12, 2018.