
An approximate KLD based experimental design for models with intractable likelihoods

Ziqiao Ao

School of Mathematical Sciences
Shanghai Jiao Tong University

Jinglai Li

School of Mathematics
University of Birmingham

Abstract

Data collection is a critical step in statistical inference and data science, and the goal of statistical experimental design (ED) is to find the data collection setup that can provide most information for the inference. In this work we consider a special type of ED problems where the likelihoods are not available in a closed form. In this case, the popular information-theoretic Kullback-Leibler divergence (KLD) based design criterion can not be used directly, as it requires to evaluate the likelihood function. To address the issue, we derive a new utility function, which is a lower bound of the original KLD utility. This lower bound is expressed in terms of the summation of two or more entropies in the data space, and thus can be evaluated efficiently via entropy estimation methods. We provide several numerical examples to demonstrate the performance of the proposed method.

1 Introduction

Collecting data is a critical step in statistical inference. In practice the data collection procedure may require considerable resources financially and in kind, and thus how to effectively allocate the limited resources in the data collection exercise becomes a question of essential importance. Statistical experimental design (ED) seeks to address the problem by developing systematic rules for allocating the resources in the data collection exercise [Ryan et al., 2016b].

In practice, the ED problems are often formulated as an optimization program, i.e., to optimize the exper-

imental conditions with respect to certain design criterion [Fedorov, 2013]. These design criteria are referred to as the utility functions, as they are meant to measure how useful the experiment outcomes, i.e., the collected data, are. It is easy to see that choosing a sensible utility functions is the key in an ED problem, and in the Bayesian inference setting, the Kullback-Leibler divergence between the prior and the posterior distributions is arguably the most popular choice for the utility function, for that it has certain theoretical merits [Paninski, 2005]. Intuitively speaking the KLD based criterion seeks to find the experimental conditions that maximize the information gain expected from conducting experiments.

As will be shown later, the use of the KLD utility requires to evaluate the likelihood function. In many real-world inference problems, e.g., the biological process models [Cook et al., 2008, Ryan et al., 2016a] the likelihood functions may be intractable, i.e., not available in a closed form, and as a result it is not possible to apply the KLD based ED directly to such problems. Considerable efforts have been devoted to developing likelihood-free ED method for such problems. Some commonly used methods rely on the ability to approximate the likelihood, or related functions, and these methods include [Overstall et al., 2018, Dehideniya et al., 2019, Kleinegesse and Gutmann, 2019]. Another popular class of likelihood free methods are based on the Approximate Bayesian Computation (ABC) [Beaumont et al., 2002] and the posterior covariance based design utility. Works in this class include [Drovandi and Pettitt, 2013, Hainy et al., 2016, Dehideniya et al., 2018], and so on. A major limitation of the posterior covariance based design utility is that it may not work when the posterior distributions deviate significantly from Gaussian.

The main goal of this work is to provide a alternative method for ED problems with intractable likelihoods, which can be used regardless whether the posterior is close to Gaussian. Specifically, we propose a utility

that is essentially a lower bound approximation of the KLD utility (namely, it is equivalent to maximizing a lower bound of the original KLD utility), and we then present a numerical scheme based on entropy estimation to evaluate this utility function without querying the likelihoods. Finally we provide numerical examples to demonstrate the performance of the proposed ED method.

The remainder of the paper is organized as follows. Section 2.1 reviews the set up of the Bayesian experimental design (BED), the KLD based design, and the methods to address intractable likelihoods. Section 3 provides our new design utility as well as an entropy estimation based numerical scheme to compute it. Finally Section 4 provides several examples to illustrate the performance of the proposed method.

2 Background

2.1 Bayesian Experimental Design

The Bayesian inference problems with controllable design parameters can be described as follows. Let $y \in \mathcal{Y}$ be the data that are observed and $\theta \in \Theta$ be the unknown parameters that we want to estimate. The relation between θ and y is characterized by the likelihood function $p(y|\theta, d)$, where d is the design parameters representing the experimental conditions that can be controlled by the users. The parameter of interest θ can be inferred from observed data y by computing the posterior distribution via Bayes' formula

$$p(\theta|y, d) = \frac{p(\theta|d)p(y|\theta, d)}{p(y|d)},$$

where $p(\theta|d)$ is the prior distribution of θ and $p(y|d)$ is the evidence of model.

As is mentioned earlier, due to time and monetary limitations, one usually can only afford to conduct a fixed number of experiments and collect their outcomes, i.e., the data. To this end, the goal of the experimental design is to identify the experimental conditions (represented by the design parameter d) which is the most useful for the inference task. Mathematically the "usefulness" of the experimental condition can be defined by a utility function: $u(d, y, \theta)$, which measures the worth of the experiment with true model parameter θ that applies design d and yields an observation y . Then one can determine the value of the design parameter d by maximizing the expected utility:

$$\max_{d \in \mathcal{D}} U(d) = \int_{\mathcal{Y}} \int_{\Theta} u(d, y, \theta) p(\theta, y|d) d\theta dy, \quad (1)$$

where \mathcal{D} is the state space of d . As is mentioned earlier, the KLD between the posterior and the prior distribu-

tions, namely,

$$u(d, y) := \int_{\Theta} p(\theta|y, d) \log \left[\frac{p(\theta|y, d)}{p(\theta|d)} \right] d\theta, \quad (2)$$

is a popular choice of the utility function for the Bayesian inference problems. One main challenge here is to solve the optimization problem in Eq. (1), and to do so we need to evaluate the objective function $U(d)$. In general the function $U(d)$ does not admit an analytical expression and needs to be evaluated numerically. First we write the function $U(d)$ as,

$$U(d) = \int_{\mathcal{Y}} \int_{\Theta} \{\log[p(y|\theta, d)] - \log[p(y|d)]\} p(y, \theta|d) dy d\theta, \quad (3)$$

It follows from Eq. (3) that $U(d)$ can be estimated via a MC simulation:

$$U(d) \approx \frac{1}{n} \sum_{i=1}^n \{\log(p(y_i|\theta_i, d)) - \log(p(y_i|d))\} \quad (4)$$

where $\{(\theta_i, y_i)\}_{i=1}^n$ are drawn from $p(y, \theta|d)$. Note however here that in usual Bayesian inference problems $p(y|d)$ is not known in advance and has to be estimated via MC as well:

$$p(y_i|d) = \int_{\Theta} p(y_i|\theta, d) p(\theta|d) d\theta \approx \frac{1}{n'} \sum_{j=1}^{n'} p(y_i|\theta_{i,j}, d), \quad (5)$$

where $\{\theta_{i,j}\}_{j=1}^{n'}$ are samples from the prior $p(\theta|d)$. By combining Eq. (4) and Eq. (5), we obtain a nested MC estimator of $U(d)$:

$$U(d) \approx \frac{1}{n} \sum_{i=1}^n \{\log(p(y_i|\theta_i, d)) - \log\left(\frac{1}{n'} \sum_{j=1}^{n'} p(y_i|\theta_{i,j}, d)\right)\}. \quad (6)$$

Once being able to evaluate the objective function $U(d)$, one can solve the optimization problem (1) by some derivative free algorithm, see e.g., [Huan and Marzouk, 2013].

2.2 BED with intractable likelihoods

We have provided a brief introduction to the BED problem in Section 2.1. It is obvious that the procedure described in Section 2.1 requires the ability to evaluate the likelihood function $p(y|\theta, d)$. which, in many complex practical problem, may not be possible. More specifically, we shall assume that, in these problems, one can draw samples from the likelihood function, but can not evaluate the likelihood function directly. To this end, several methods have been proposed to conduct the likelihood-free experimental design, and most of them are based on the ABC method for sampling the posterior. Though these methods are

different technically, the main ideas behind them are quite similar: one should seek the design parameter that yields the posterior distribution with the minimal ‘‘uncertainty’’, which is measured by certain quantities associated with the posterior covariance. The posterior distribution here is computed by the ABC method.

We give a brief description of one of this type methods. For example, one can chose the utility function to be the inverse of the determinant of the posterior covariance matrix:

$$u(d, y) = \frac{1}{\det(\text{cov}(\theta|y, d))}, \quad (7)$$

which is numerically evaluated by,

$$u(d, y) \approx \frac{1}{\det(\hat{C}(\theta_1, \dots, \theta_n))}, \quad (8)$$

where $\{\theta_i\}_{i=1}^n$ is an ensemble drawn from the posterior distribution $p(\theta|y, d)$ and \hat{C} is the sample covariance of it. Obviously to evaluate such a $u(d, y)$ one need to draw samples from the posterior distribution $p(\theta|y, d)$. Using the ABC methods [Beaumont et al., 2002], one can draw samples from the posterior distribution without evaluating the likelihood. This way, one can evaluate the utility function $u(d, y)$ for any given pair of (d, y) . Finally, the expected utility $U(d)$ is then computed via MC:

$$U(d) = \frac{1}{n} \sum_{i=1}^n u(d, y_i), \quad (9)$$

where $\{y_i\}_{i=1}^n$ are drawn from the joint distribution $p(y, \theta)$. Following [Drovandi and Pettitt, 2013], we refer to this approach as the *D-posterior precision* method. Other utility function associated with the posterior covariance can also be used, e.g. the inverse of the trace of the posterior covariance. By using the ABC method and utility function based on the posterior covariance, one can conduct the experimental design without evaluating the likelihood function. However, a major limitation of this type of methods is that the use of such utility may be inappropriate if the posterior distributions deviate significantly from Gaussian (e.g., multimodal distributions). as in that case the uncertainty of the posterior may not be well quantified by its variance.

The aforementioned KLD utility can alleviate this limitation and be used for any posterior distribution. Unfortunately, as is pointed out at the beginning of the section, computing the KLD utility is extremely challenging when the likelihood is intractable. Several works such as [Overstall et al., 2018], [Kleinegesse and Gutmann, 2019] propose to address the issue by approximating or numerically

estimating the likelihood or related functions, and accurate estimation or approximation of these functions may be either computationally intensive or assumes the likelihood to be of certain specific form. An alternative solution is to evaluate the expected KLD utility by entropy estimation. To this end, [Terejanu et al., 2012] evaluates the expected utility by estimating the mutual information in the joint space of θ and y , which is impractical when the unknown θ is high dimensional. In this work, we adopt the entropy estimation framework, and in the next section we propose a new utility function which only requires to estimate a small number of entropies in the output space.

3 The entropy estimation based BED method

In this section we introduce the approximate KLD utility and its numerical implementation.

3.1 A lower bound approximation of the expected KLD utility

By some elementary calculus we can rewrite Eq. (3) as,

$$U(d) = -E_{\theta}[H(p(y|\theta, d))] + H(p(y|d)), \quad (10)$$

where $H(p(\cdot))$ is defined as the entropy of distribution $p(\cdot)$:

$$H[p(y)] = - \int \log[p(y)]p(y)dy. \quad (11)$$

A very intuitive idea to evaluate $U(d)$ is to apply an MC estimation to the first term in Eq. (10):

$$U(d) \approx - \sum_{i=1}^n H(p(y|\theta_i, d)) + H(p(y|d)), \quad (12)$$

where $\theta^1 \dots \theta^n$ are drawn from the prior $p(\theta)$. One then can use some entropy estimation approach to compute $H(p(y|\theta, d))$ and $H(p(y))$. Even though the procedure described above can evaluate $U(d)$ with intractable likelihood function, it may not be practical as it requires to perform entropy estimation $n + 1$ times. In reality, the entropy estimation is computationally intensive, and n , the number of MC samples are usually large (say, 10^4 or larger), which may render the total computational cost prohibitive. Given that we want to maximize $U(d)$, an intuitive solution is to construct a lower bound approximation to $U(d)$, and maximize this lower bound instead. Our construction of the lower bound is based on Theorem 3.1.

Theorem 3.1 *Suppose that θ is a random variable defined on state space Θ , with probability density $p(\theta)$.*

For any given $\theta \in \Theta$, let y and y' be two random variables that are independent conditional on θ , and both follow the same distribution $p(y|\theta)$. Now define $z = y - y'$, and we then have,

$$E_{\theta}[H(p(y|\theta))] \leq H(E_{\theta}[p(z|\theta)]) - \frac{\dim(y)}{2} \log 2,$$

where $\dim(y)$ is the dimensionality of y .

The proof of the Theorem 3.1 is based on the entropy power inequality [Cover and Thomas, 2012] and entropy's concavity property [Cover and Thomas, 2012], and a complete proof is provided in the Supplementary Material. Applying Theorem 3.1 to Eq. (10) one obtains a lower bound of $U(d)$:

$$U_L(d) = -H(E_{\theta}[p(z|\theta, d)]) + \frac{\dim(y)}{2} \log 2 + H(p(y|d)) \quad (13)$$

Note here that to evaluate this approximation one only need to perform entropy estimation twice. Moreover, this lower bound approximation can be further refined, based on the following corollary:

Corollary 3.2 *Suppose $p(\theta)$, $p(y|\theta)$, and $p(z|\theta)$ are defined as is in Theorem 3.1, and $p(\theta)$ admits the form of,*

$$p(\theta) = \sum_{l=1}^L \omega_l f_l(\theta),$$

where $\omega_l \geq 0$ for $l = 1 \dots L$, $\sum_{l=1}^L \omega_l = 1$, and $f_l(\theta)$ are density functions. Then

$$\begin{aligned} E_{\theta}[H(p(y|\theta))] &\leq \sum_{l=1}^L \omega_l H(E_{\theta \sim f_l}[p(z|\theta)]) - \frac{\dim(y)}{2} \log 2 \\ &\leq H(E_{\theta}[p(z|\theta)]) - \frac{\dim(y)}{2} \log 2. \end{aligned}$$

The proof of the Corollary is also provided in the Supplementary Material. Now we discuss how to design a lower bound based on Corollary 3.2. First we divide the observation space Y into a fixed number of disjoint partitions: $\{Y_l\}_{l=1}^L$, such that $\cup_{l=1}^L Y_l = Y$ and $Y_l \cap Y_{l'} = \emptyset$ if $l \neq l'$. It then follows that the prior distribution can be written as

$$p(\theta) = \sum_{l=1}^L \omega_l f_l(\theta), \quad (14)$$

where $f_l(\theta) = p(\theta|y \in Y_l)$ and $w_l = p(y \in Y_l)$ for $l = 1 \dots L$. In turn we obtain another lower bound approximation of $U(d)$:

$$\begin{aligned} U_L(d) &= - \sum_{l=1}^L \omega_l H(E_{\theta \sim f_l}[p(z|\theta)]) \\ &\quad + \frac{\dim(y)}{2} \log 2 + H(p(y)), \quad (15) \end{aligned}$$

which is tighter than that in Eq. (13). In what follows we shall refer to $U_L(d)$ in Eq. (15) as the expected lower-bound (LB)-KLD utility with partition and $U_L(d)$ in Eq. (13) as the expected LB-KLD utility without partition. In our method, we choose to use LB-KLD with partition. Next we shall discuss how to numerically evaluate this new expected utility function.

3.2 Evaluating the lower bound approximation

First note that the proposed lower bound approximation $U_L(d)$ in Eq. (15) is actually the sum of several entropies. Thus an effective entropy estimation method is required to compute this approximation. In this work we adopt the Nearest Neighbor (NN) based entropy estimator provided in [Kraskov et al., 2004], while noting that other choices are also available [Beirlant et al., 1997]. Using the method, we can estimate the entropy from a given set of samples.

We summarize the procedure of estimating the lower bound approximation $U_L(d)$ in Algorithm 1. In addition to the design parameter value d , the algorithm users need to specify two algorithm parameters: the number of samples generated n , and the number of partitions of the state space L . In this algorithm the function EntEst[.] takes a set of samples as its input and outputs an entropy of these samples estimated using the method in [Kraskov et al., 2004]. A key step in the algorithm is to decompose the prior distribution $p(\theta|d)$ into a number of mixture components as is in Eq. (14). As is mentioned earlier, the mixture representation of $p(\theta|d)$ is constructed by partitioning the output space \mathcal{Y} , and numerically the partition is achieved by clustering the samples in the output space into a fixed number of groups. Specifically, we use a constrained k-means method [Berkhin, 2006] to cluster the samples of into L groups subject to the constraint that the size of any group is larger than a fixed value n_{\min} . Once the clusters of the output samples are determined, the mixture of $p(y|\theta)$ is obtained. As one can see that, in the algorithm, we do not need an explicit form of the mixture, and rather we only need the clustering of the samples drawn from the prior distribution.

4 Numerical Examples

In this section, we demonstrate the performance of the proposed design method with three examples. In all these examples, comparisons are made between the proposed method and the D-posterior precision method (with the posterior samples generated via ABC). The implementation details such as the algo-

Algorithm 1 LB-KLD Estimator

Input: d, n, L
Output: \hat{U}_L

- 1: **for** $i = 1$ to n **do**
- 2: Generate $\theta_i \sim p(\theta|d)$ and $y_i^* \sim p(y|\theta_i, d)$
- 3: **end for**
- 4: $\hat{H}^* = \text{EntEst}[\{y_i^*\}_{i=1}^n]$
- 5: $[\tilde{\Theta}, N] = \text{Partition}[\{(\theta_i, y_i^*)\}_{i=1}^n, L]$
- 6: **for** $l = 1$ to L **do**
- 7: $\omega_l = \frac{N(l)}{n}$
- 8: **for** $j = 1$ to $j = N(l)$ **do**
- 9: Generate $y_j, y'_j \sim p(y|x_j \sim \Theta_l, d)$ // x_j represents the j th element of Θ_l
- 10: Compute $z_j = y_j - y'_j$
- 11: **end for**
- 12: $\hat{H}_l = \text{EntEst}[\{z_j\}_{j=1}^{N(l)}]$
- 13: **end for**
- 14: $\hat{U}_L = -\sum_{l=1}^L \omega_l \hat{H}_l + \frac{\dim(y)}{2} \log 2 + \hat{H}^*$
- 15: **return** \hat{U}_L

Algorithm 2 $[\Theta, N] = \text{Partition}[\{(\theta_i, y_i^*)\}_{i=1}^n, L]$

 User-specified parameters: n_{\min}

- 1: Use the constrained k-means method to cluster $\{y_i^*\}_{i=1}^n$ into L groups subject to the constraint that the size of any group is no smaller than n_{\min} ;
 - 2: Cluster $\{\theta_i\}_{i=1}^n$ into L groups $\Theta_1, \dots, \Theta_k$ according to the clustering results of $\{y_i^*\}_{i=1}^n$;
 - 3: $\tilde{\Theta} = \{\Theta_l\}_{l=1}^L$;
 - 4: **for** $l = 1$ to L **do**
 - 5: $N(l) = \text{NumOf}(\Theta_l)$;
 - 6: **end for**
-

rithm parameters are provided in the supplemental material. We also note that, the main purpose of the examples is to compare the performance of the ED methods, and detailed comparisons in terms of both performance and computational cost will be reported elsewhere.

4.1 A mathematical example

To illustrate the limitation of the ABC based method, we first consider a toy problem with strongly non-Gaussian posterior. Specifically, consider the following generative model

$$y = G(\theta, d)(1 + \epsilon_1) + \epsilon_2, \quad G(\theta, d) = \frac{1}{B(2, d)}\theta(1 - \theta)^{d-1},$$

where $B(\cdot, \cdot)$ is the beta function, $\epsilon_1 \sim N(0, 0.05^2)$ and $\epsilon_2 \sim N(0, 0.05^2)$. In this example, the likelihood

is actually available,

$$p(y|\theta, d) = N(G(\theta, d), 0.05^2(1 + G^2(\theta, d))).$$

The prior is assumed to be uniformly distributed on the interval $[0, 1]$. The design variable is chosen from $[2, 100]$. The main purpose of this example is two-fold: first since the likelihood is available in this problem, we can accurately evaluate the expected utility, and validate our approximation against it; second, the posterior distribution of this problem is strongly non-Gaussian, and we thus can show that the KLD and the ABC methods yield very different design in this case.




Figure 1: The KLD expected utility plotted against the design parameter d , compared to that based on the D-posterior, computed with ABC.

First we estimate the KLD based expected utility function with the nested MC method, which is regarded as the exact value of the KLD. We then compute the lower bound approximations of the expected utility with and without the partition refinement. We plot the expected utility computed by all the three methods as a function of the design parameter d in Fig. 1. As one can see here that, the approximation computed with partition agree very well with the exact value of the expected utility, while that without partition admits obvious discrepancy from the exact value of the expected utility. Nevertheless, we can see from the figure that the optimal solutions predicted by all three methods are largely the same $d = 5$. As a comparison, in Fig. 1 we plot the expected D-posterior precision utility as a function of d . Interestingly, in this case the expected utility is roughly an increasing function of d , and as a result the optimal solution is achieved at $d = 100$, which is the upper bound of d .

To further compare the two methods, we conduct numerical experiments under the two experimental conditions: $d = 5$ and $d = 100$. We generate data and perform the Bayesian inference for two cases: the true value is $\theta = 0.5$ and the true value is $\theta = 0.8$. We

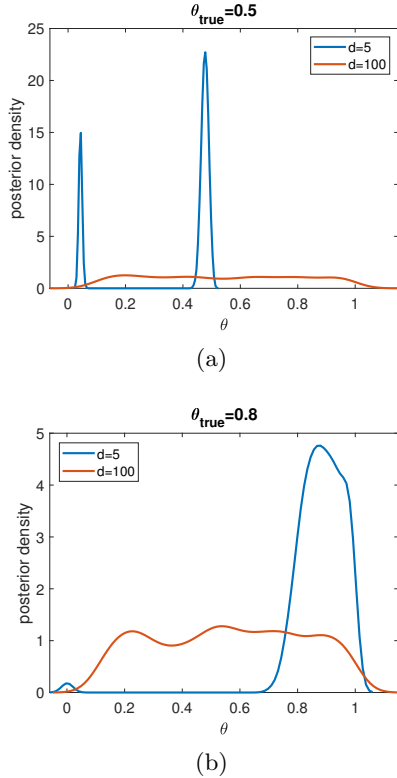


Figure 2: The posterior distributions for $\theta_{true} = 0.5$ (a) and $\theta_{true} = 0.8$ (b), obtained under the two experimental conditions $d = 5$ and $d = 100$.

show the obtained posterior distributions in Fig. 2. One can see here that the posterior distributions in this problem are greatly different from Gaussian and those obtained under $d = 5$ are clearly bimodal with one mode concentrated around the ground truth. On the other hand, clearly the posteriors obtained under the condition $d = 100$ look substantially less informative than those under $d = 5$, suggesting that the KLD utility is more effective for problems with strongly non-Gaussian posteriors.

4.2 Ricker Model

Our second example is a ecological model describing the evolution of population size over time. The model assumes that the unobservable population size N_t follows the scaled Ricker map [Turchin, 2003]:

$$N_{t+1} = rN_t e^{-N_t + e_t}, \quad t = 1, \dots, T$$

where $e_t \sim N(0, \sigma^2)$ are independent process noise and r is a parameter related to the growth rate. The observation Y_t is Poisson distributed as:

$$Y_t \sim \text{Poisson}(\phi N_t),$$

where ϕ is a scale parameter. We assume $\theta = (\log r, \phi, \sigma)$ are the three parameters to be inferred

and the number of observations is $T = 50$. Following [Dutta et al., 2016], the prior distribution is given by

$$\log r \sim U(3, 5), \quad \phi \sim U(5, 15), \quad \sigma \sim U(0, 0.6).$$

For convenience's sake, one does not use the observations y_1, \dots, y_t directly, and instead we select a number of summary statistics of the observations and use them as the data for the inference. Following [Wood, 2010], we set 13 summary statistics of $\{y_1, \dots, y_T\}$: the average population and number of zeros observed over the given time, the autocovariances from lag0 to lag5, the coefficients α_0, α_1 and α_2 of the quadratic regression $y_{t+1} = \alpha_2(y_{t+1} - y_t)^2 + \alpha_1(y_{t+1} - y_t) + \alpha_0 + \epsilon_t$, and the autoregression coefficients β_0 and β_1 based on the regression $y_{t+1}^{0.3} = \beta_0 y_t^{0.3} + \beta_1 y_t^{0.6} + \epsilon_t$. For convenience's sake, we index these statistics as is in Table 1. The goal here is to find a combination of two statistics in the 13 which can provide most information. The likelihood function here is obviously not available in a closed form, and it is an often used example for likelihood-free inference [Price et al., 2018, Fearnhead and Prangle, 2012].

Table 1: numbered statistics

No.	statistics	No.	statistics
1	average population	8	autocov lag5
2	zeros observed	9	α_0
3	autocov lag0	10	α_1
4	autocov lag1	11	α_2
5	autocov lag2	12	β_0
6	autocov lag3	13	β_1
7	autocov lag4		

We compute the expected LB-KLD utility at all the scenarios of the combination and show the results in Fig. 3 (a), which indicate that the optimal combination should be that of statistics (1) and (2). As a comparison, we also compute the expected D-posterior precision utility, also for all scenarios, shown in Fig. 3 (b), and with this utility function we obtain a different optimal combination, i.e., statistics (2) and (3). Thus, in this example, the two methods also result in different experimental conditions. To compare the performance of the two combinations, we conduct the following experiments. First we randomly generate 1000 pairs of truth θ_{true} and simulated data y from the distribution $p(y, \theta)$. With each pair of (θ_{truth}, y) , we conduct a Bayesian inference and compute the posterior distribution using statistics (1) and (2), and using (2) and (3) respectively. Since the likelihood is not available, all the posteriors are computed with the ABC method. In Fig 4 we show the scatter plots of the

posterior means versus the true values of the 1000 trials for all the three parameters. Fig. 4 (a) shows the results from statistics 1&2, which is the optimal condition determined by the LB-KLD method, and Fig. 4 (b) shows those from statistics 2&3, the optimal solution by the D-posterior precision method. The solid lines which represents that the posterior value is equal to the truth provides a guideline of the inference results, and intuitively speaking, the points being distributed closer to the line indicates better inference results. In this respect, we can see that parameter σ is much more difficult to infer than the other two parameters, in both designs. However, the figures show that design obtained by the LB-KLD utility seems to lead to better inference results for $\log r$ and ϕ . The advantage of the proposed method can also be quantitatively demonstrated by calculating the mean square error of the posterior means in the 1000 trials. While the MSEs for σ are similar: 0.080 for LB-KLD and 0.079 for D-posterior precision, the design obtained by our method results in much lower MSEs for $\log r$ and ϕ : 0.01 (LB-KLD) versus 0.02 (D-posterior) for $\log r$ and 0.0046 (LB-KLD) versus 0.0073 (D-posterior) for ϕ respectively.

4.3 Aphid Model

The last problem we consider is a stochastic model describing the growth of aphid population [Matis et al., 2007]. The purpose of this example is to illustrate that when the posterior distribution is not too far from Gaussian, the proposed method identify similar designs as the D-posterior precision approach. Let $N(t)$ and $C(t)$ denote the current live size and the accumulative size of aphid population. The population is assumed to grow with a rate of $\lambda N(t)$ and die with a rate of $\mu N(t)C(t)$ at any time t . Therefore, the probability that a birth or a death occurs in the next infinitesimal time period Δ_t is

$$\begin{aligned} Pr\{N(t+\Delta_t)=n+1, C(t+\Delta_t)=c+1|N(t)=n, C(t)=c\} \\ = \lambda n \Delta_t + o(\Delta_t), \\ Pr\{N(t+\Delta_t)=n-1, C(t+\Delta_t)=c|N(t)=n, C(t)=c\} \\ = \mu n c \Delta_t + o(\Delta_t), \end{aligned}$$

where λ and μ are the birth and death parameters to be determined. The prior distributions for λ and μ follow [Gillespie and Boys, 2019], where

$$\begin{pmatrix} \lambda \\ \mu \end{pmatrix} = N \left[\begin{pmatrix} 0.246 \\ 0.000136 \end{pmatrix}, \begin{pmatrix} 0.0079^2 & 5.8 \times 10^{-8} \\ 5.8 \times 10^{-8} & 0.00002^2 \end{pmatrix} \right].$$

And the initial aphid level is set as $N(0) = C(0) = 28$. The goal here is to specify the sampling times $D = (t_1, \dots, t_k)$ so as to accurately estimate the model parameters. This problem also does not admit a

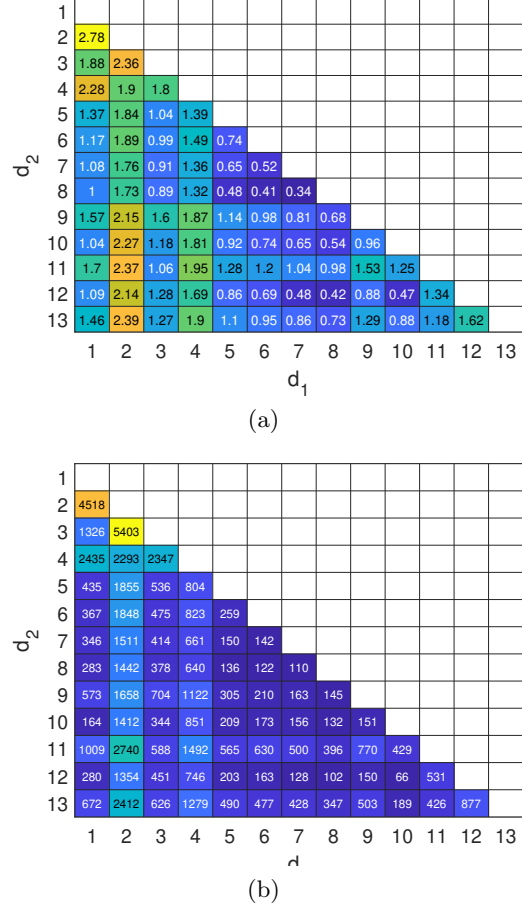


Figure 3: (a): the expected LB-KLD utility at all the scenarios. (b): the expected D-posterior precision utility at all the scenarios.

tractable likelihood function and so it can not be solved with the standard KLD based method.

We consider the observation times design in time interval $[0, 50]$ which is equally discretized into 5000 grids for the purpose of searching for the optimal designs. We then compute the optimal designs for $k = 1, 2, 3, 4$ using the LB-KLD and the D-posterior precision methods, and the optimal designs are provided in Table 2 One can see from the table that in this example, the methods actually identify very close experimental conditions. To further analyze the methods, we conduct a Bayesian inference for the situation where the truth is located at the mean of the prior, and the simulated data is collected at 4 time locations: $(13.8, 19.1, 24.5, 30.6)$, which are the optimal choice selected by the LB-KLD method. We compute the posterior with ABC and plot the posterior distributions in Fig. 5, and also shown in the figure are the Gaussian fits of the posterior distributions. The figures show that the posterior distributions in this problem are very close to Gaussian, which provides a good ex-

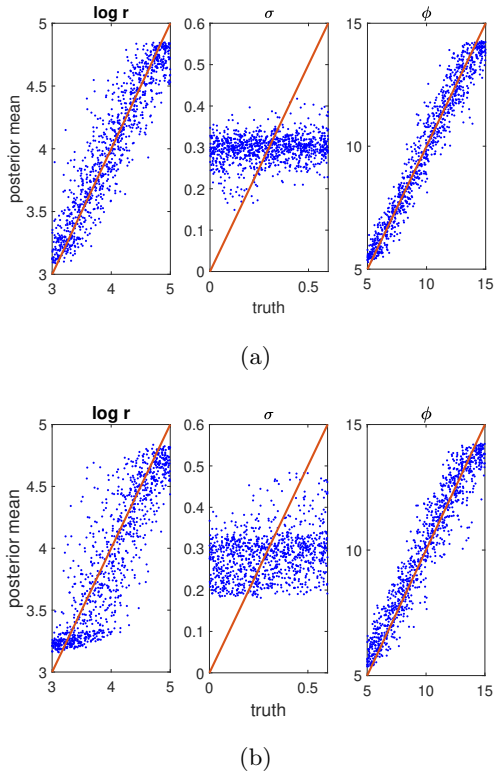


Figure 4: The scatter plot of the posterior means against the true values for $d = (1, 2)$ (the top) and $d = (2, 3)$ (the bottom).

planation that the optimal designs obtained by both methods are rather close to each other.

Table 2: The optimal design results for the Aphid Model.

method	LB-KLD	D-posterior
1d	(21)	(21)
2d	(17, 28)	(18, 27)
3d	(15.7, 22.7, 32.0)	(16.8, 21.9, 29.1)
4d	(13.8, 19.1, 24.5, 30.6)	(15.8, 20.4, 25.2, 30.5)

5 Discussion

In this paper we have presented a Bayesian experimental design method for stochastic models with intractable likelihoods, and the method can be applied to problems where the posteriors are strongly non-Gaussian. Specifically we propose a new utility function, which is a lower bound approximation to the often used KLD utility, and we provide an entropy estimation based method to estimate the expected utility. Using numerical examples, we demonstrate that, the method performs well regardless whether the posterior is close to Gaussian, while the D-posterior precision

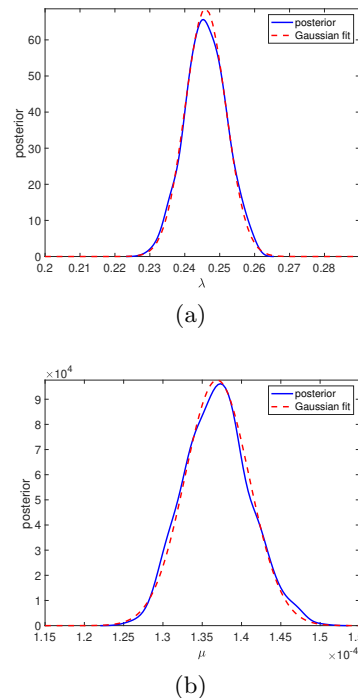


Figure 5: (a): the posterior of λ (solid) and its Gaussian fit (dashed). (b): the posterior of μ (solid) and its Gaussian fit (dashed).

method yields less effective designs when the posteriors differs strongly from Gaussian. We believe the method can be found useful in a large range of real world BED problems where the likelihood functions are intractable.

A number of improvements and extensions of the proposed LB-KLD method are possible. First as is mentioned earlier, in this work we only compare the performance against the D-posterior method, and a more comprehensive comparison with many other aforementioned methods e.g. [Kleingesse and Gutmann, 2019], in terms of both effectiveness and efficiency, will be conducted in a future work. Secondly, as the LB-KLD is essentially a lower bound approximation of the KLD utility, we anticipate that in certain circumstances, the method may not yield the same results as the KLD, in which case, the approximation method is deemed unsuitable. Thus in the future we hope to understand the applicability as well as the limitations of the method. Last but not least, the prior partition procedure is an essential component of the proposed method. In this work we provide an empirical partition strategy based on clustering the output samples, and an important question here is that what is the optimal partition strategy (including the number of partitions L) in theory? We plan to investigate these issues in future studies.

Acknowledgements

The work was partially supported by NSFC under grant number 11301337.

References

- [Beaumont et al., 2002] Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035.
- [Beirlant et al., 1997] Beirlant, J., Dudewicz, E. J., Györfi, L., and Van der Meulen, E. C. (1997). Non-parametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6(1):17–39.
- [Berkhin, 2006] Berkhin, P. (2006). A survey of clustering data mining techniques. In *Grouping multidimensional data*, pages 25–71. Springer.
- [Cook et al., 2008] Cook, A. R., Gibson, G. J., and Gilligan, C. A. (2008). Optimal observation times in experimental epidemic processes. *Biometrics*, 64(3):860–868.
- [Cover and Thomas, 2012] Cover, T. M. and Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.
- [Dehideniya et al., 2019] Dehideniya, M., Overstall, A. M., Drovandi, C. C., and McGree, J. M. (2019). A synthetic likelihood-based laplace approximation for efficient design of biological processes. *arXiv preprint arXiv:1903.04168*.
- [Dehideniya et al., 2018] Dehideniya, M. B., Drovandi, C. C., and McGree, J. M. (2018). Optimal bayesian design for discriminating between models with intractable likelihoods in epidemiology. *Computational Statistics & Data Analysis*, 124:277–297.
- [Drovandi and Pettitt, 2013] Drovandi, C. C. and Pettitt, A. N. (2013). Bayesian experimental design for models with intractable likelihoods. *Biometrics*, 69(4):937–948.
- [Dutta et al., 2016] Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. (2016). Likelihood-free inference by ratio estimation. *arXiv preprint arXiv:1611.10242*.
- [Fearnhead and Prangle, 2012] Fearnhead, P. and Prangle, D. (2012). Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):419–474.
- [Fedorov, 2013] Fedorov, V. V. (2013). *Theory of optimal experiments*. Elsevier.
- [Gillespie and Boys, 2019] Gillespie, C. S. and Boys, R. J. (2019). Efficient construction of bayes optimal designs for stochastic process models. *Statistics and Computing*, pages 1–10.
- [Hainy et al., 2016] Hainy, M., Müller, W. G., and Wagner, H. (2016). Likelihood-free simulation-based optimal design with an application to spatial extremes. *Stochastic Environmental Research and Risk Assessment*, 30(2):481–492.
- [Huan and Marzouk, 2013] Huan, X. and Marzouk, Y. M. (2013). Simulation-based optimal bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317.
- [Kleinegesse and Gutmann, 2019] Kleinegesse, S. and Gutmann, M. U. (2019). Efficient bayesian experimental design for implicit models. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 476–485. PMLR.
- [Kraskov et al., 2004] Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Physical review E*, 69(6):066138.
- [Matis et al., 2007] Matis, J. H., Kiffe, T. R., Matis, T. I., and Stevenson, D. E. (2007). Stochastic modeling of aphid population growth with nonlinear, power-law dynamics. *Mathematical biosciences*, 208(2):469–494.
- [Overstall et al., 2018] Overstall, A., McGree, J., et al. (2018). Bayesian design of experiments for intractable likelihood models using coupled auxiliary models and multivariate emulation. *Bayesian Analysis*.
- [Paninski, 2005] Paninski, L. (2005). Asymptotic theory of information-theoretic experimental design. *Neural Computation*, 17(7):1480–1507.
- [Price et al., 2018] Price, L. F., Drovandi, C. C., Lee, A., and Nott, D. J. (2018). Bayesian synthetic likelihood. *Journal of Computational and Graphical Statistics*, 27(1):1–11.
- [Ryan et al., 2016a] Ryan, C. M., Drovandi, C. C., Pettitt, A. N., et al. (2016a). Optimal bayesian experimental design for models with intractable likelihoods using indirect inference applied to biological process models. *Bayesian Analysis*, 11(3):857–883.
- [Ryan et al., 2016b] Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2016b). A review

of modern computational algorithms for bayesian optimal design. *International Statistical Review*, 84(1):128–154.

[Terejanu et al., 2012] Terejanu, G., Upadhyay, R. R., and Miki, K. (2012). Bayesian experimental design for the active nitridation of graphite by atomic nitrogen. *Experimental Thermal and Fluid Science*, 36:178–193.

[Turchin, 2003] Turchin, P. (2003). *Complex population dynamics: a theoretical/empirical synthesis*, volume 35. Princeton university press.

[Wood, 2010] Wood, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102.