

# Approximate Bayesian inference from noisy likelihoods with Gaussian process emulated MCMC

**Marko Järvenpää**

*Department of Biostatistics, University of Oslo, Norway*

M.J.JARVENPAA@MEDISIN.UIO.NO

**Jukka Corander**

*Department of Biostatistics, University of Oslo, Norway*

*Department of Mathematics and Statistics, University of Helsinki, Finland*

*Wellcome Sanger Institute, United Kingdom*

JUKKA.CORANDER@MEDISIN.UIO.NO

**Editor:** John Cunningham

## Abstract

We present a framework for approximate Bayesian inference intended for a situation where only a limited number of noisy log-likelihood evaluations can be obtained due to constraints on the available computational budget, which is becoming increasingly common for expensive simulator-based models. We model the log-likelihood function using a Gaussian process (GP) and our main methodological innovation is to apply this model to emulate the progression that an exact Metropolis-Hastings (MH) sampler would take if it was applicable. Informative log-likelihood evaluation locations are selected using a sequential experimental design strategy until the MH accept/reject decisions are performed with sufficient level of accuracy based on a prespecified error tolerance criterion. The resulting approximate sampler is conceptually simple and shown to be sample-efficient. It is also more robust compared with earlier “Bayesian optimisation-like” methods tailored for approximate Bayesian inference, which generally assume a global surrogate model across the parameter space that can be challenging to fit well. We discuss some theoretical aspects and various interpretations of the resulting approximate MH sampler, and demonstrate its benefits in the context of Bayesian and generalised Bayesian likelihood-free inference for simulator-based statistical models.

**Keywords:** approximate Bayesian inference, Markov chain Monte Carlo, Gaussian process, likelihood-free inference, sequential experimental design

## 1. Introduction

Standard computational methods for Bayesian inference, such as Markov chain Monte Carlo (MCMC), require typically a large number of likelihood evaluations and are consequently poorly suited for situations where the likelihood function is expensive to evaluate. Additional challenge to MCMC is brought by noisy likelihood evaluations which occur in the context of likelihood-free inference (LFI), where the intractable likelihood function is itself estimated from forward simulations of the statistical model, see Marin et al. (2012); Sisson et al. (2019); Cranmer et al. (2020). In particular, in the synthetic likelihood (SL) method (Wood, 2010; Price et al., 2018), typically hundreds or even thousands of repeated simulations are needed to approximate the likelihood function sufficiently well at each considered point in the parameter space, which is costly and still results in noisy estimates of the un-

derlying SL function. Other examples of expensive likelihood evaluations include parameter estimation in complex ordinary/partial differential equations with tractable noise models (e.g. Cleary et al. (2021); Paun and Husmeier (2022)) and Bayesian machine learning with large data sets (e.g. Korattikara et al. (2014); Bardenet et al. (2017)).

Different methods based on GP surrogate models have been earlier proposed to reduce the computational cost of approximate Bayesian inference for simulator-based models. In particular, recently Järvenpää et al. (2021) proposed a framework for situations where only a very limited number (on the order of  $10^3$ ) of noisy log-likelihood evaluations can be obtained, referred to as “Bayesian LFI” (BLFI) in this paper. BLFI is related to so-called probabilistic numerics methods (Hennig et al., 2015; Cockayne et al., 2019; Hennig et al., 2022) and the key idea is to frame the computation of the posterior density itself as a Bayesian inference task, in the same spirit as global, derivative-free optimisation is framed as an inference task in the acclaimed Bayesian optimisation methodology. In BLFI the log-likelihood function is modelled with a GP which is further used to derive an estimator for the posterior density. In this context, *active learning* is the essential strategy to maximise the information obtainable from a limited number of log-likelihood evaluations.

BLFI and other related approaches such as the operationally similar *Bayesian optimisation for likelihood-free inference* (BOLFI) framework by Gutmann and Corander (2016) use a global GP surrogate model over the whole parameter space to achieve sample-efficiency. While this strategy has been shown to work well in many cases, some difficulties may also arise under various circumstances. Firstly, the global modelling approach is sensible when the target density is expected to be highly multimodal, but often the posterior is actually close to a Gaussian, is unimodal or at least its highest density region corresponds to a relatively small region of the parameter space. Consequently, the global surrogate model and the associated acquisition functions may become inefficient, in particular when dimensionality of the parameter space increases. Secondly, the likelihood function can also behave irregularly, as is often the case with nonlinear dynamic models used in ecology and epidemiology (Fasiolo et al., 2016), or be arbitrarily small near the parameter boundaries which results practical difficulties with global GP model fitting.

The currently used active learning strategies for B(O)LFI suffer from a number of issues. First, they all require global optimisation of an *acquisition function* at each iteration of the inference algorithm. In particular, the one-step ahead optimal Bayesian experimental design strategies developed by Järvenpää et al. (2021) are rather costly, and although this is only a minor concern when the likelihood evaluations are truly expensive, it still markedly complicates the inference workflow. The cheaper acquisition functions based on uncertainty sampling also considered by Järvenpää et al. (2021), do not come with theoretical guarantees and exhibited unreliable performance in numerical experiments. The ones borrowed from Bayesian optimisation literature (see Gutmann and Corander (2016)) are not theoretically fully sound as they are not explicitly designed to estimate the posterior distribution (Kandasamy et al., 2017; Järvenpää et al., 2019, 2021), and even though they have been demonstrated to work well in many cases, they tend to produce quite a few redundant evaluations near the boundaries as observed e.g. in Järvenpää et al. (2019); Picchini et al. (2022).

In this paper we develop and analyse a new inference framework GP-MH that directly combines MH sampling with the benefits of the B(O)LFI approach and alleviates the afore-

mentioned difficulties. In particular, the resulting GP-MH algorithm models and explores the parameter space in a more local fashion by emulating the progression of an exact, but as such directly inapplicable MH sampler. This allows to either avoid or robustly manage the problematic likelihood evaluations especially near the parameter boundaries. While possible information available about the shape or location of the posterior density region associated with substantial probability mass (based on pilot runs, expert knowledge or iterative model building), is difficult to code directly into the GP model, it can be more pragmatically captured in an informed initialisation scheme or the proposal density used by the sampler. Sequential Bayesian experimental design strategies are developed here to aid the algorithm selecting informative evaluation locations and they provide fairly similar sample-efficiency as the existing B(O)LFI methods but are more computationally efficient and better interpretable. The resulting GP-MH implementation, similarly to the current B(O)LFI methods, is mainly intended for models with low- to moderate-dimensional parameter space and requires assumptions regarding the smoothness of the underlying log-likelihood function, both of which are generally assumed to be met in likelihood-free inference for simulator-based statistical models.

The remaining paper is organised as follows. We first provide background on MH sampling, BLFI and previous methods in statistics and machine learning literature that use GPs for more efficient Bayesian inference. We then develop our framework and derive several sequential experimental design strategies for it. In Section 4 we observe that some common MH samplers are special cases of GP-MH and discuss how GP-MH can be interpreted as a special case of BLFI, which motivates an alternative implementation called MH-BLFI, or as a heuristic, yet tractable estimate to a conceptual ‘‘Bayesian’’ MH sampler. Section 4.4 outlines some theoretical aspects of GP-MH. In section 5 we investigate the proposed implementations numerically, first with synthetically constructed target densities and then with realistic simulator-based models using SL and generalised Bayesian updating in the spirit of Bissiri et al. (2016); Pacchiardi and Dutta (2021). Summary and discussion about future research directions conclude the paper.

## 2. Approximate Bayesian inference

We denote the observed data as  $\mathbf{x}_o \in \mathbb{R}^d$ , the unknown parameters of a statistical model as  $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p$  and the associated likelihood function as  $\pi(\mathbf{x}_o | \boldsymbol{\theta})$ . We focus on continuous parameter spaces but most of the analysis extends to the case where some components of  $\boldsymbol{\theta}$  are discrete. The prior density is denoted by  $\pi(\boldsymbol{\theta})$  and is assumed tractable to evaluate directly. We denote the posterior distribution as  $\pi(\boldsymbol{\theta} | \mathbf{x}_o) = \pi(\boldsymbol{\theta})\pi(\mathbf{x}_o | \boldsymbol{\theta})/\pi(\mathbf{x}_o)$ , where  $\pi(\mathbf{x}_o) = \int_{\Theta} \pi(\boldsymbol{\theta}')\pi(\mathbf{x}_o | \boldsymbol{\theta}') d\boldsymbol{\theta}'$  is the marginal likelihood. Throughout this paper the main goal will be on sampling from  $\pi(\boldsymbol{\theta} | \mathbf{x}_o)$  or computing posterior expectations of the form

$$\bar{h} \triangleq \mathbb{E}_{\boldsymbol{\theta} | \mathbf{x}_o} h(\boldsymbol{\theta}) = \int_{\Theta} h(\boldsymbol{\theta})\pi(\boldsymbol{\theta} | \mathbf{x}_o) d\boldsymbol{\theta} = \frac{\int_{\Theta} h(\boldsymbol{\theta})\pi(\boldsymbol{\theta})\pi(\mathbf{x}_o | \boldsymbol{\theta}) d\boldsymbol{\theta}}{\int_{\Theta} \pi(\boldsymbol{\theta}')\pi(\mathbf{x}_o | \boldsymbol{\theta}') d\boldsymbol{\theta}'}, \quad (1)$$

where  $h : \Theta \rightarrow \mathbb{R}$  is some known, tractable function that is cheap to evaluate.

Metropolis-Hastings sampler (Hastings, 1970) is widely used for Monte Carlo integration. Algorithm 1 shows in a compact form how MH is used to draw samples from  $\pi(\boldsymbol{\theta} | \mathbf{x}_o)$ . Under certain technical conditions, this MH sampler produces a Markov chain whose stationary

distribution is the posterior  $\pi(\boldsymbol{\theta} | \mathbf{x}_o)$ . The algorithm starts from an initial point  $\boldsymbol{\theta}^{(0)}$ . At each iteration  $i$  a new parameter denoted by  $\boldsymbol{\theta}'^{(i)}$  is drawn from the proposal density  $q(\boldsymbol{\theta}'^{(i)} | \boldsymbol{\theta}^{(i-1)})$  and is then accepted with probability  $\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$ , where

$$\alpha(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq \min\{1, \gamma(\boldsymbol{\theta}, \boldsymbol{\theta}')\}, \quad \gamma(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq \frac{\pi(\boldsymbol{\theta}')\pi(\mathbf{x}_o | \boldsymbol{\theta}')q(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta})\pi(\mathbf{x}_o | \boldsymbol{\theta})q(\boldsymbol{\theta}' | \boldsymbol{\theta})}, \quad (2)$$

and otherwise the current point  $\boldsymbol{\theta}^{(i-1)}$  is kept. The initial samples (e.g. the first half) are often discarded as “burn-in”. The remaining samples, here denoted as  $\boldsymbol{\theta}^{(0)}, \dots, \boldsymbol{\theta}^{(n)}$ , are approximately distributed as  $\pi(\boldsymbol{\theta} | \mathbf{x}_o)$  and can be used to estimate (1) as

$$\bar{h} \approx \hat{h}_{n+1} \triangleq \frac{1}{n+1} \sum_{i=0}^n h(\boldsymbol{\theta}^{(i)}). \quad (3)$$

See e.g. Robert and Casella (2004) for a more detailed treatment of MCMC methods.

---

**Algorithm 1** Metropolis-Hastings (MH) sampler for  $\pi(\boldsymbol{\theta} | \mathbf{x}_o)$

---

**Input:** Prior  $\pi(\boldsymbol{\theta})$ , likelihood  $\pi(\mathbf{x}_o | \boldsymbol{\theta})$ , initial point  $\boldsymbol{\theta}^{(0)}$ , proposal  $q(\boldsymbol{\theta}' | \boldsymbol{\theta})$ , no. samples  $i_{\text{MH}}$

**Output:** Samples  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(i_{\text{MH}})}$

- 1: **for**  $i = 1 : i_{\text{MH}}$  **do**
  - 2:     Draw  $\boldsymbol{\theta}'^{(i)} \sim q(\cdot | \boldsymbol{\theta}^{(i-1)})$  and  $u^{(i)} \sim \mathcal{U}([0, 1])$
  - 3:     Set  $\boldsymbol{\theta}^{(i)} \leftarrow \boldsymbol{\theta}'^{(i)} \mathbb{1}_{\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)}) \geq u^{(i)}} + \boldsymbol{\theta}^{(i-1)} \mathbb{1}_{\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)}) < u^{(i)}}$
  - 4: **end for**
- 

## 2.1 Bayesian inference with noisy and expensive likelihoods

In the BLFI framework the log-likelihood function, from now on denoted by  $f : \Theta \rightarrow \mathbb{R}$  so that  $f(\boldsymbol{\theta}) \triangleq \log \pi(\mathbf{x}_o | \boldsymbol{\theta})$  where the dependence on the fixed data  $\mathbf{x}_o$  is suppressed for brevity, is itself treated as an unknown random function to be estimated using another level of Bayesian inference. A GP prior is assigned on  $f$  and a Gaussian noise model is assumed for the log-likelihood evaluations. Some  $t$  parameter values are then chosen to evaluate the log-likelihood  $f$  which results in the data  $\mathcal{D}_t$ . Then  $f | \mathcal{D}_t$  follows a GP which further induces a log-GP posterior for the *unnormalised* posterior

$$\tilde{\pi}_f(\boldsymbol{\theta}) \triangleq \pi(\boldsymbol{\theta})e^{f(\boldsymbol{\theta})}. \quad (4)$$

An estimate for (4) is obtained analytically by using the properties of GP models. This GP-based estimate defines an approximation to the exact unnormalised posterior  $\tilde{\pi}(\boldsymbol{\theta} | \mathbf{x}_o) \triangleq \pi(\boldsymbol{\theta})\pi(\mathbf{x}_o | \boldsymbol{\theta})$  and is plugged in to a standard MCMC algorithm. This last step does not require further log-likelihood evaluations as it is solely based on the GP posterior. Finally, the resulting samples can be used to approximate  $\bar{h}$  as usual, via (3).

The new evaluation location is chosen as the global optimum of an acquisition function at each iteration of BLFI. In the case of sequential Bayesian experimental design, the acquisition function is defined as an expected loss function where the loss quantifies the uncertainty associated with the unnormalised posterior (4) and the expectation is taken

---

**Algorithm 2** General form of B(O)LFI implementation

---

**Input:** Prior  $\pi(\boldsymbol{\theta})$ , Bayesian model (e.g. GP) for the target log-likelihood  $f(\boldsymbol{\theta})$ , no. initial evaluations  $t_{\text{init}}$ , no. total evaluations  $t_{\text{max}}$ , no. MCMC samples  $s_{\text{MCMC}}$

**Output:** Samples  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(s_{\text{MCMC}})}$

- 1: **for**  $j = 1 : t_{\text{init}}$  **do** ▷ Collect data for the initial model fitting.
  - 2:     Select initial evaluation location  $\boldsymbol{\theta}_j$  ▷ E.g. draw  $\boldsymbol{\theta}_j \sim_{\text{i.i.d.}} \mathcal{U}(\Theta)$  or  $\boldsymbol{\theta}_j \sim_{\text{i.i.d.}} \pi(\boldsymbol{\theta})$ .
  - 3:     Set  $y_j \leftarrow$  log-likelihood evaluation at  $\boldsymbol{\theta}_j$
  - 4: **end for**
  - 5: Set  $t \leftarrow t_{\text{init}}$  and  $\mathcal{D}_t \leftarrow \{(y_j, \boldsymbol{\theta}_j)\}_{j=1}^t$
  - 6: Fit the Bayesian model for  $f$  using  $\mathcal{D}_t$
  - 7: **for**  $t = t_{\text{init}} + 1 : t_{\text{max}}$  **do** ▷ Collect data for improving the model.
  - 8:     Select informative evaluation location  $\boldsymbol{\theta}^*$  ▷ By optimising an acquisition function.
  - 9:     Set  $y^* \leftarrow$  log-likelihood evaluation at  $\boldsymbol{\theta}^*$
  - 10:    Set  $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup \{(y^*, \boldsymbol{\theta}^*)\}$
  - 11:    Refit Bayesian model for  $f$  using  $\mathcal{D}_t$
  - 12: **end for**
  - 13: Sample  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(s_{\text{MCMC}})}$  from a model-based estimate of (4) with MCMC
- 

with respect to a hypothetical future evaluation based on the current GP model. A general implementation of B(O)LFI in a concise form is outlined as Algorithm 2.

The approximation error can also be assessed probabilistically, at least in principle. For example, the expectation  $\bar{h}$  in (1), which is here interpreted as a functional of  $f$  and rewritten as

$$\bar{h}_f \triangleq \int_{\Theta} h(\boldsymbol{\theta}) \pi_f(\boldsymbol{\theta}) \, d\boldsymbol{\theta}, \quad \pi_f(\boldsymbol{\theta}) \triangleq \frac{\pi(\boldsymbol{\theta}) e^{f(\boldsymbol{\theta})}}{\int_{\Theta} \pi(\boldsymbol{\theta}') e^{f(\boldsymbol{\theta}')} \, d\boldsymbol{\theta}'},$$

follows a posterior distribution induced by the GP of  $f | \mathcal{D}_t$  via the non-linear mapping  $f \mapsto \bar{h}_f$ . The resulting density  $\pi(\bar{h} | \mathcal{D}_t)$  can be assessed numerically (see Järvenpää et al. (2020)) though this approach is itself approximate and works only in low dimensions.

Aided by the more efficient use of a limited budget of log-likelihood evaluations, BLFI typically needs only a few hundred evaluations for reasonable posterior approximations, which is significantly less than required by pseudo-marginal/noisy MH. BLFI is conceptually similar to Bayesian optimisation (Hennig and Schuler, 2012; Frazier, 2018; Garnett, 2023), adaptive warped Bayesian quadrature (Osborne et al., 2012; Gunter et al., 2014; Chai and Garnett, 2019) and GP-based noisy level set estimation (Bect et al., 2012; Lyu et al., 2021) developed for other related numerical tasks involving expensive functions.

## 2.2 Other related literature

Accelerating MCMC by using GPs or other surrogate models, also called emulators or metamodels, has been widely considered in literature, see Llorente et al. (2021) for a recent survey. Notably, Rasmussen (2003); Christen and Fox (2005); Bliznyuk et al. (2008); Fielding et al. (2011); Conrad et al. (2016); Zhang et al. (2017); Sherlock et al. (2017); Zhang and Taflanidis (2019) develop asymptotically exact MCMC algorithms mainly in the context of exact but costly likelihood evaluations often resulting from complex ODE or PDE systems

with tractable observation models. Sometimes derivative evaluations are also available to aid GP fitting (Lan et al., 2016; Paun and Husmeier, 2022). Different from these studies, we instead focus on expensive stochastic models whose likelihood function is approximated using forward simulations. We also aim for the best possible sample-efficiency instead of merely improving over standard MCMC, while accepting some bias due to the use of the GP surrogate. Related techniques that assume expensive likelihood evaluations but which are not directly based on MCMC include Kandasamy et al. (2017); Wang and Li (2018); Acerbi (2018); Alawieh et al. (2020). These methods use global GP modelling and we expect them to consequently suffer from similar modelling challenges as B(O)LFI. Moreover, the numerical experiments by Järvenpää et al. (2021); Acerbi (2020) suggest that the active learning strategies used in these related approaches are not robust in the noisy setting.

Bayesian inference using MH sampling in the case of “tall data” (Korattikara et al., 2014; Angelino et al., 2016; Bardenet et al., 2017; Zhang et al., 2020) is another related and likewise challenging task. While the underlying statistical model is tractable and typically relatively cheap, the very large number of data points causes likelihood evaluations to be costly. A key idea in these cases is to use unbiased log-likelihood evaluations obtained by subsampling the data points in the MH accept/reject test. Although the existing methods are better tailored for this specific problem, our proposed approach could in principle also be applied there.

In addition to B(O)LFI, other inference frameworks based on GP surrogate modelling have been proposed for LFI. Our approach most closely resembles the GPS-ABC algorithm by Meeds and Welling (2014) where a related approximate MH framework is considered. However, a major difference is that in GPS-ABC individual summary statistics are modelled with independent GPs in the context of Approximate Bayesian Computation (ABC) while we model the log-likelihood with GP (not necessarily in the ABC scenario). In addition, we provide substantially more comprehensive analysis of the main idea and extend it in various ways in our setting. In Wilkinson (2014) the difficulties with global GP modelling are partially eluded by classifying problematic parameter regions as implausible at each “wave” of their algorithm and fitting the GP only to its complement. However, this approach is in practice difficult to automatise. Finally, GP-accelerated MCMC methods in the context of noisy log-likelihood evaluations have been considered by Drovandi et al. (2018); Wiqvist et al. (2018) while variational inference was used by Acerbi (2020). However, these methods are quite convoluted featuring multiple stages and are not designed for maximal sample-efficiency.

### 2.3 Gaussian process emulated MH with noisy likelihood evaluations

We now develop our framework for emulating the progression of an exact MH when only a limited number of noisy log-likelihood evaluations are accessible. We model the log-likelihood using a probabilistic surrogate model and explicitly treat  $\gamma(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$ —and consequently also  $\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$ —as random variables. In Section 2.4 we discuss how the MH accept/reject decisions, that control the progression of the MH algorithm, are made in the presence of uncertainty of the exact value of  $\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$  in the light of Bayesian decision theory. Then we present a particular GP surrogate model for the log-likelihood

function (Section 2.5), combine this model with the preceding theory (Section 2.6) and finally form a practical implementation of the framework (Section 2.7).

## 2.4 Uncertainty in the MH acceptance ratio

Let us revisit the MH sampler in Algorithm 1. An essential initial observation is that  $\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$  in line 3 can equivalently be replaced by the slightly simpler quantity  $\gamma(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$ . Whether to accept or reject a proposed  $\boldsymbol{\theta}'^{(i)}$  at iteration  $i$ , when the current point is  $\boldsymbol{\theta}^{(i-1)}$  and when there is uncertainty about the corresponding likelihood values, is treated as a problem of Bayesian decision theory. Potential previous or future decisions are not taken into account for simplicity. Let  $\hat{\gamma} = \hat{\gamma}(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$  be an estimator for the random variable  $\gamma = \gamma(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)})$  for making the decision. We consider a loss function

$$l_u(\gamma, \hat{\gamma}) \triangleq \mathbb{1}_{\gamma < u, \hat{\gamma} \geq u} + \mathbb{1}_{\gamma \geq u, \hat{\gamma} < u}$$

with a fixed  $u = u^{(i)} \in [0, 1]$ . The loss is 1 if we choose  $\hat{\gamma} \geq u$  while in reality  $\gamma < u$  or if we choose  $\hat{\gamma} < u$  while  $\gamma \geq u$ , and 0 otherwise. Both type of errors are hence considered equally undesirable. The expected loss is then

$$\begin{aligned} \mathbb{E}_\gamma l_u(\gamma, \hat{\gamma}) &= \int_{\mathbb{R}} (\mathbb{1}_{\gamma < u} \mathbb{1}_{\hat{\gamma} \geq u} + \mathbb{1}_{\gamma \geq u} \mathbb{1}_{\hat{\gamma} < u}) dF_\gamma(\gamma) \\ &= \mathbb{1}_{\hat{\gamma} \geq u} \int_{\mathbb{R}} \mathbb{1}_{\gamma < u} dF_\gamma(\gamma) + \mathbb{1}_{\hat{\gamma} < u} \int_{\mathbb{R}} \mathbb{1}_{\gamma \geq u} dF_\gamma(\gamma) \\ &= \mathbb{P}(\gamma < u \mid u) \mathbb{1}_{\hat{\gamma} \geq u} + \mathbb{P}(\gamma \geq u \mid u) \mathbb{1}_{\hat{\gamma} < u}, \end{aligned} \tag{5}$$

where  $F_\gamma(\gamma)$  is the cumulative distribution function (CDF) of  $\gamma$ . We also define an alternative loss function  $l(\gamma, \hat{\gamma}) \triangleq \int_0^1 l_u(\gamma, \hat{\gamma}) du$ . Using Fubini's theorem we obtain the expected loss  $\mathbb{E}_\gamma l(\gamma, \hat{\gamma}) = \int_0^1 \mathbb{E}_\gamma l_u(\gamma, \hat{\gamma}) du$ , whose integrand is given by (5).

Similarly to Meeds and Welling (2014), we may also consider the probability of making an error in the MH acceptance test. This is done either conditionally on  $u$  so that

$$\begin{aligned} \mathcal{E}_{u, \hat{\gamma}} &\triangleq \mathbb{P}(\text{“Incorrect accept/reject decision”} \mid \hat{\gamma}, u) \\ &= \mathbb{P}(\{\gamma < u, \hat{\gamma} \geq u\} \cup \{\gamma \geq u, \hat{\gamma} < u\} \mid \hat{\gamma}, u) \\ &= \mathbb{P}(\gamma < u, \hat{\gamma} \geq u \mid \hat{\gamma}, u) + \mathbb{P}(\gamma \geq u, \hat{\gamma} < u \mid \hat{\gamma}, u) \\ &= \mathbb{P}(\gamma < u \mid u) \mathbb{1}_{\hat{\gamma} \geq u} + \mathbb{P}(\gamma \geq u \mid u) \mathbb{1}_{\hat{\gamma} < u}, \end{aligned} \tag{6}$$

or unconditionally by averaging over  $u \sim \mathcal{U}([0, 1])$  so that

$$E_{\hat{\gamma}} \triangleq \int_0^1 \mathcal{E}_{u, \hat{\gamma}} \mathcal{U}(u \mid [0, 1]) du = \int_0^1 \mathcal{E}_{u, \hat{\gamma}} du. \tag{7}$$

We see that (6), which we refer simply as the *conditional error* from now on, coincides with the expected loss (5). Similarly, the *unconditional error* (7) equals  $\mathbb{E}_\gamma l(\gamma, \hat{\gamma})$ . An optimal estimator  $\hat{\gamma}$  for making the accept/reject decision minimises the expected loss. Recall that the *median* of a real-valued random variable  $z$ , which we denote by  $\text{med}(z)$ , is defined as any value  $m \in \mathbb{R}$  satisfying  $\mathbb{P}(z \leq m) \geq 1/2$  and  $\mathbb{P}(z \geq m) \geq 1/2$ . The median always exists but may not be unique.

**Proposition 1** *Suppose  $\gamma$  is a real-valued random variable. Then the choice  $\hat{\gamma} = \text{med}(\gamma)$  (where  $\text{med}(\gamma)$  can be any of its median values) minimises the unconditional error  $E_{\hat{\gamma}}$  and also the conditional error  $\mathcal{E}_{u, \hat{\gamma}}$  for each fixed  $u \in [0, 1]$ .*

The proof for this and some other theoretical results are given in Appendix B. From now on we use  $\hat{\gamma}$  exclusively to denote this optimal estimator. It follows that the optimal decision is to choose the most probable action given  $u$  because

$$\begin{aligned} \mathcal{E}_{u, \hat{\gamma}} &= \mathbb{P}(\gamma < u | u) \mathbf{1}_{\text{med}(\gamma) \geq u} + \mathbb{P}(\gamma \geq u | u) \mathbf{1}_{\text{med}(\gamma) < u} \\ &= \min\{\mathbb{P}(\gamma < u | u), \mathbb{P}(\gamma \geq u | u)\}. \end{aligned} \quad (8)$$

In the following sections we show that, unlike in the different surrogate modelling scenario of Meeds and Welling (2014), analytical formulas for the key quantities above can be obtained when the log-likelihood follows GP posterior.

## 2.5 GP model for the log-likelihood

We present a GP model which is similar to the one used by Järvenpää et al. (2021) and is suitable for low-dimensional parameters ( $p < 10$ ). While we use this model in this paper, other choices may be more appropriate in some other settings. We first assume

$$y_j = f(\boldsymbol{\theta}_j) + \epsilon_j, \quad \epsilon_j \sim \mathcal{N}(0, \sigma_n^2(\boldsymbol{\theta}_j)), \quad j = 1, \dots, t, \quad (9)$$

where  $y_j \in \mathbb{R}$  denotes a noisy evaluation of the log-likelihood function  $f$  at some parameter  $\boldsymbol{\theta}_j \in \Theta$  and  $\sigma_n^2 : \Theta \rightarrow \mathbb{R}_+$  is the noise variance. The Gaussian noise assumption (9) is further discussed in Appendix E.1.

We then place the following hierarchical GP prior for  $f$ :

$$f | \boldsymbol{\beta} \sim \mathcal{GP}(m_0(\boldsymbol{\theta}), k_\phi(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad m_0(\boldsymbol{\theta}) = \sum_{i=1}^q \beta_i h_i(\boldsymbol{\theta}), \quad \boldsymbol{\beta} \sim \mathcal{N}(\mathbf{b}, \mathbf{B}), \quad (10)$$

where  $k_\phi : \Theta \times \Theta \rightarrow \mathbb{R}$  is a covariance (kernel) function with hyperparameters  $\phi$  and  $h_i : \Theta \rightarrow \mathbb{R}$  denote fixed basis functions. The covariance function encodes the smoothness assumption of the likelihood function whereas possible prior assumptions on its shape can be encoded by specifying suitable basis functions. In our analysis we assume that  $\phi$  and  $\sigma_n^2(\boldsymbol{\theta})$  for each  $\boldsymbol{\theta}$  are known though in practice their values are obtained using point estimation. We omit  $\phi$  from our notation for brevity.

As in O'Hagan and Kingman (1978); Rasmussen and Williams (2006), we integrate out  $\boldsymbol{\beta}$  in (10). Given evaluations  $\mathcal{D}_t \triangleq \{(y_j, \boldsymbol{\theta}_j)\}_{j=1}^t$ , the posterior of  $f$  can be shown to be  $f | \mathcal{D}_t \sim \mathcal{GP}(m_t(\boldsymbol{\theta}), c_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$ , where

$$\begin{aligned} m_t(\boldsymbol{\theta}) &\triangleq \mathbf{k}_t(\boldsymbol{\theta}) \mathbf{K}_t^{-1} \mathbf{y}_t + \mathbf{R}_t^\top(\boldsymbol{\theta}) \bar{\boldsymbol{\beta}}_t, \\ c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') &\triangleq k(\boldsymbol{\theta}, \boldsymbol{\theta}') - \mathbf{k}_t(\boldsymbol{\theta}) \mathbf{K}_t^{-1} \mathbf{k}_t^\top(\boldsymbol{\theta}') + \mathbf{R}_t^\top(\boldsymbol{\theta}) [\mathbf{B}^{-1} + \mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{H}_t^\top]^{-1} \mathbf{R}_t(\boldsymbol{\theta}'), \end{aligned}$$

with  $[\mathbf{K}_t]_{ij} \triangleq k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) + \mathbf{1}_{i=j} \sigma_n^2(\boldsymbol{\theta}_i)$  for  $i, j = 1, \dots, t$ ,  $\mathbf{k}_t(\boldsymbol{\theta}) \triangleq (k(\boldsymbol{\theta}, \boldsymbol{\theta}_1), \dots, k(\boldsymbol{\theta}, \boldsymbol{\theta}_t))$ ,  $\bar{\boldsymbol{\beta}}_t \triangleq [\mathbf{B}^{-1} + \mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{H}_t^\top]^{-1} (\mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{y}_t + \mathbf{B}^{-1} \mathbf{b})$  and  $\mathbf{R}_t(\boldsymbol{\theta}) \triangleq \mathbf{H}(\boldsymbol{\theta}) - \mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{k}_t^\top(\boldsymbol{\theta})$ . The columns of  $\mathbf{H}_t \in \mathbb{R}^{q \times t}$  consist of basis function values evaluated at  $\boldsymbol{\theta}_{1:t} = [\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_t] \in \mathbb{R}^{p \times t}$



and  $\mathbf{H}(\boldsymbol{\theta})$  is the corresponding  $q \times 1$  vector at  $\boldsymbol{\theta}$ . We also have  $\mathbf{y}_t = (y_1, \dots, y_t)^\top$  and we additionally denote the GP variance function as  $s_t^2(\boldsymbol{\theta}) \triangleq c_t(\boldsymbol{\theta}, \boldsymbol{\theta})$ . See Rasmussen and Williams (2006) for further details on GP regression and Appendix E.2 for remarks on modelling log-likelihood functions.

## 2.6 Uncertainty in the MH acceptance ratio based on GP surrogate

We apply the analysis of Section 2.4 on handling the uncertainty in the MH accept/reject test when the log-likelihood function follows a GP posterior as in Section 2.5. Here  $\boldsymbol{\theta}$  denotes the current point at an arbitrary iteration of the MH sampler and  $\boldsymbol{\theta}'$  is the corresponding proposal generated from  $q(\boldsymbol{\theta}' | \boldsymbol{\theta})$ . We have

$$\begin{bmatrix} f(\boldsymbol{\theta}) \\ f(\boldsymbol{\theta}') \end{bmatrix} | \mathcal{D}_t \sim \mathcal{N}_2 \left( \begin{bmatrix} m_t(\boldsymbol{\theta}) \\ m_t(\boldsymbol{\theta}') \end{bmatrix}, \begin{bmatrix} s_t^2(\boldsymbol{\theta}) & c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') \\ c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') & s_t^2(\boldsymbol{\theta}') \end{bmatrix} \right),$$

which further implies

$$f(\boldsymbol{\theta}') - f(\boldsymbol{\theta}) | \mathcal{D}_t \sim \mathcal{N}(m_t(\boldsymbol{\theta}') - m_t(\boldsymbol{\theta}), s_t^2(\boldsymbol{\theta}') + s_t^2(\boldsymbol{\theta}) - 2c_t(\boldsymbol{\theta}, \boldsymbol{\theta}')). \quad (11)$$

Using (2), which we can rewrite as

$$\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq \frac{\pi(\boldsymbol{\theta}') e^{f(\boldsymbol{\theta}')} q(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta}) e^{f(\boldsymbol{\theta})} q(\boldsymbol{\theta}' | \boldsymbol{\theta})} = \frac{\pi(\boldsymbol{\theta}') q(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta}) q(\boldsymbol{\theta}' | \boldsymbol{\theta})} e^{f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})},$$

and (11), it follows that  $\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')$  given evaluations  $\mathcal{D}_t$  follows log-Normal distribution:

$$\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') | \mathcal{D}_t \sim \log \mathcal{N}(\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}'), \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad (12)$$

$$\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq m_t(\boldsymbol{\theta}') - m_t(\boldsymbol{\theta}) + \log \left( \frac{\pi(\boldsymbol{\theta}') q(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta}) q(\boldsymbol{\theta}' | \boldsymbol{\theta})} \right), \quad (13)$$

$$\sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq s_t^2(\boldsymbol{\theta}') + s_t^2(\boldsymbol{\theta}) - 2c_t(\boldsymbol{\theta}, \boldsymbol{\theta}'). \quad (14)$$

Furthermore,  $\alpha_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq \min\{1, \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')\}$  given  $\mathcal{D}_t$  follows a mixture density consisting of a log-Normal density in  $[0, 1)$  and a point mass at 1. Its CDF is  $F_{\alpha_f(\boldsymbol{\theta}, \boldsymbol{\theta}') | \mathcal{D}_t}(a) = \Phi((\log(a) - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')) / \sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')) \mathbf{1}_{a < 1} + \mathbf{1}_{a \geq 1}$  for  $a > 0$  and  $F_{\alpha_f(\boldsymbol{\theta}, \boldsymbol{\theta}') | \mathcal{D}_t}(a) = 0$  for  $a \leq 0$ , where  $\Phi(\cdot)$  is the CDF of the standard Gaussian distribution.

Given the GP posterior of  $f | \mathcal{D}_t$  and the optimal estimator

$$\hat{\gamma} = \hat{\gamma}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \text{med}_{f | \mathcal{D}_t} \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') = e^{\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')} \quad (15)$$

for  $\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')$ , the conditional and unconditional errors defined in Section 2.4 are

$$\mathcal{E}_{t,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \Phi \left( -\frac{|\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \log(u)|}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')} \right), \quad (16)$$

$$E_{t,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \int_0^1 \Phi \left( -\frac{|\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \log(u)|}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')} \right) du \quad (17)$$

$$= \begin{cases} \Phi(-\mu_t/\sigma_t) - e^{\mu_t + \sigma_t^2/2} \Phi(-(\mu_t + \sigma_t^2)/\sigma_t) & \text{if } \mu_t \geq 0, \\ \Phi(\mu_t/\sigma_t) + e^{\mu_t + \sigma_t^2/2} [\Phi(-(\mu_t + \sigma_t^2)/\sigma_t) - 2\Phi(-\sigma_t)] & \text{if } \mu_t < 0, \end{cases} \quad (18)$$




Figure 1: The coloured solid lines show the conditional error  $\mathcal{E}_{t,u,\hat{\gamma}}$  in (16) with some choices of  $u$  and the dashed black line is the unconditional error  $E_{t,\hat{\gamma}}$  in (18). With each realisation of  $u$ ,  $\mu_t$  can be chosen so that  $\mathcal{E}_{t,u,\hat{\gamma}}$  equals its maximal value  $\Phi(0) = 1/2$ . When the GP posterior becomes more accurate in the sense that  $\sigma_t$  decreases, the intervals containing such  $\mu_t$  values that lead to unconfident MH accept/reject decision become narrower. Importantly, negligible (un)conditional error results when  $|\mu_t|$  is large as compared to  $\sigma_t$  in which case the accept/reject decision can be done with high confidence although the corresponding likelihood ratio is not necessarily known accurately.

respectively. Note that we sometimes shorten  $\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  as  $\mu_t$  and similarly  $\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  as  $\sigma_t$ . Equations (16) and (18) are derived in Appendix B and illustrated in Figure 1.

We can also define

$$\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') \triangleq \mathbb{1}_{\alpha_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq u} = \mathbb{1}_{\log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \log(u)}$$

so that  $\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 1$  if  $\boldsymbol{\theta}'$  is to be accepted and  $\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 0$  otherwise (for a given  $u \in [0, 1]$  and  $f$ ). We define  $\tilde{u} \triangleq \log(u)$  and see immediately that  $\mathbb{P}_{f|\mathcal{D}_t}(\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 0) = \Phi((\tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))/\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$  and  $\mathbb{E}_{f|\mathcal{D}_t} \kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \mathbb{P}_{f|\mathcal{D}_t}(\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 1) = \Phi((\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u})/\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$ . It further holds that

$$\mathbb{V}_{f|\mathcal{D}_t} \kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \Phi((\tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))/\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))\Phi((\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u})/\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')). \quad (19)$$

We also see that

$$\mathcal{E}_{t,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \min\{\mathbb{P}(\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 0 | \mathcal{D}_t), \mathbb{P}(\kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 1 | \mathcal{D}_t)\},$$

which shows that the most probable decision given the GP posterior is here made, as also implied by (8). Using the equations above, the fact  $\Phi(z) = 1 - \Phi(-z)$  and the inequality  $\min\{x, 1 - x\} \leq \sqrt{x(1 - x)}$  for  $x \in [0, 1]$ , we also see that the conditional error  $\mathcal{E}_{t,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is upper bounded by  $(\mathbb{V}_{f|\mathcal{D}_t} \kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}'))^{1/2}$ .

## 2.7 GP-MH implementation

We combine the preceding analysis and GP model to form Algorithm 3. At each iteration  $i$  of the resulting GP-MH algorithm, the proposal  $\boldsymbol{\theta}'^{(i)}$  is either accepted or rejected based on the GP posterior conditioned on the  $t$  evaluations in  $\mathcal{D}_t$  collected that far. The decision is made in a greedy optimal manner based on the results in Section 2.4 and 2.6. In a

similar spirit to Meeds and Welling (2014); Korattikara et al. (2014), new log-likelihood evaluations are acquired (lines 10-11) until the estimated probability of making an incorrect MH accept/reject decision is smaller than a pre-specified tolerance parameter  $\varepsilon$  (line 9). In most iterations, this is achieved without new log-likelihood evaluations which facilitates computational savings. Details on the selection of the evaluation locations (line 10) are given in Section 3. The GP model is updated after each new evaluation (line 13). The outputted samples are finally used to approximate the posterior or some posterior expectations of interest via (3). Note that two distinct parameter sets are maintained in the algorithm:  $\boldsymbol{\theta}_j$  in  $\mathcal{D}_t$  denote the evaluation locations for the GP fit whereas  $\boldsymbol{\theta}^{(i)}$  denote the resulting approximate MH samples. Handling of possible problematic log-likelihood evaluations on lines 3 and 11 are described in Appendix E.2.

---

**Algorithm 3** Approximate GP-emulated MH (GP-MH)

---

**Input:** Prior  $\pi(\boldsymbol{\theta})$ , GP model for  $f$ , no. initial evaluations  $t_{\text{init}}$ , error tolerance  $\varepsilon$ , initial point  $\boldsymbol{\theta}^{(0)}$ , proposal  $q(\boldsymbol{\theta}' | \boldsymbol{\theta})$ , no. MH samples  $i_{\text{MH}}$   
**Output:** Approximate MH samples  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(i_{\text{MH}})}$

- 1: **for**  $j = 1 : t_{\text{init}}$  **do** ▷ Obtain evaluations for the initial GP fitting.
- 2:     Sample  $\boldsymbol{\theta}_j \sim_{\text{i.i.d.}} q(\cdot | \boldsymbol{\theta}^{(0)})$  ▷ Other initial points can also be used.
- 3:     Set  $y_j \leftarrow$  log-likelihood evaluation at  $\boldsymbol{\theta}_j$
- 4: **end for**
- 5: Set  $t \leftarrow t_{\text{init}}$  and  $\mathcal{D}_t \leftarrow \{(y_j, \boldsymbol{\theta}_j)\}_{j=1}^t$
- 6: Fit GP using  $\mathcal{D}_t$
- 7: **for**  $i = 1 : i_{\text{MH}}$  **do**
- 8:     Sample  $\boldsymbol{\theta}'^{(i)} \sim q(\cdot | \boldsymbol{\theta}^{(i-1)})$  and  $u^{(i)} \sim \mathcal{U}([0, 1])$
- 9:     **while**  $E_{t, \hat{\gamma}}(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}'^{(i)}) > \varepsilon$  **do** ▷ Alternatively, use  $\mathcal{E}_{t, u^{(i)}, \hat{\gamma}}$ .
- 10:         Obtain  $\boldsymbol{\theta}^*$  as a solution to (25) ▷ See Section 3.
- 11:         Set  $y^* \leftarrow$  log-likelihood evaluation at  $\boldsymbol{\theta}^*$
- 12:         Set  $t \leftarrow t + 1$  and  $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup \{(y^*, \boldsymbol{\theta}^*)\}$
- 13:         Refit GP using  $\mathcal{D}_t$
- 14:     **end while**
- 15:     Set  $\boldsymbol{\theta}^{(i)} \leftarrow \boldsymbol{\theta}'^{(i)} \mathbb{1}_{\hat{\gamma} \geq u^{(i)}} + \boldsymbol{\theta}^{(i-1)} \mathbb{1}_{\hat{\gamma} < u^{(i)}}$  ▷ Accept/reject  $\boldsymbol{\theta}'^{(i)}$ ;  $\hat{\gamma}$  computed using (15).
- 16: **end for**

---

In this paper we consider a random-walk Metropolis version of GP-MH with  $q(\boldsymbol{\theta}' | \boldsymbol{\theta}) = \mathcal{N}_p(\boldsymbol{\theta}' | \boldsymbol{\theta}, \boldsymbol{\Sigma})$ . It is often difficult to select a suitable proposal covariance  $\boldsymbol{\Sigma}$  *a priori*. In our implementation we specify an initial covariance matrix  $\boldsymbol{\Sigma}_0$  and update it based on the obtained samples as in the adaptive Metropolis algorithm by Haario et al. (2001). We use the initial proposal density  $\mathcal{N}_p(\boldsymbol{\theta}' | \boldsymbol{\theta}, \boldsymbol{\Sigma}_0)$  also to obtain evaluations around  $\boldsymbol{\theta}^{(0)}$  for initial GP fitting (lines 1-4). Contrary to a typical MCMC use case, possible poor mixing is not a major concern in our setting where the parameter space is low-dimensional, time spent on evaluating the log-likelihood dominates and the MH accept/reject decision is based solely on the GP on most iterations. Finding a good initial location and proposal covariance is still beneficial and pilot runs may be needed.

### 3. One-step ahead optimal evaluation locations

We choose the evaluation locations in a one-step ahead optimal manner in the sense of Bayesian experimental design theory, see e.g. Garnett (2023, Chapter 5) for background. In this so-called “myopic” strategy only the effect of the next log-likelihood evaluation (or a batch of evaluations) is taken into account while potential additional future evaluations needed to make the (un)conditional error eventually smaller than  $\varepsilon$  or possible later MH transitions are not. This common strategy substantially simplifies the computations. We denote a collection of candidate evaluation locations as  $\boldsymbol{\theta}^* \in \mathbb{R}^{p \times b}$  and the corresponding log-likelihood evaluations as  $\mathbf{y}^* \in \mathbb{R}^b$ , where  $b \geq 1$  is the batch size, that is, the number of simultaneous evaluations. We also denote  $\mathcal{D}^* \triangleq \{(y_j^*, \boldsymbol{\theta}_j^*)\}_{j=1}^b$ . Algorithm 3 is stated for the sequential case  $b = 1$  but we present the theory for the more general batch case  $b \geq 1$  as this comes with little additional difficulty.

Specifically, such evaluation location(s)  $\boldsymbol{\theta}^*$  are treated optimal that minimise the expected loss where the loss function quantifies the uncertainty associated with the current MH accept/reject decision. Suitable choices include the conditional and unconditional errors. The expectation is taken with respect to the predictive density of the future log-likelihood evaluation(s) at  $\boldsymbol{\theta}^*$  given by  $\pi(\mathbf{y}^* | \boldsymbol{\theta}^*, \mathcal{D}_t) = \mathcal{N}_b(\mathbf{y}^* | m_t(\boldsymbol{\theta}^*), c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \boldsymbol{\Lambda}^*)$  where  $\boldsymbol{\Lambda}^* \triangleq \text{diag}(\sigma_n^2(\boldsymbol{\theta}_1^*), \dots, \sigma_n^2(\boldsymbol{\theta}_b^*))$ .

**Proposition 2** *Suppose the GP model in Section 2.5 holds and consider the above set-up. The expected conditional error  $L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$ , the expected unconditional error  $L_t^E(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  and the expected variance of  $\kappa_{u,f}$  denoted by  $L_t^V(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  are then given by*

$$L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) \triangleq \mathbb{E}_{\mathbf{y}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} \mathcal{E}_{t+b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 2T \left( \frac{\tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')} , \frac{\sqrt{\sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)}}{\xi_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)} \right),$$

$$L_t^E(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) \triangleq \mathbb{E}_{\mathbf{y}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} E_{t+b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \int_0^1 L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) du, \quad (20)$$

$$\begin{aligned} L_t^{V,u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) &\triangleq \mathbb{E}_{\mathbf{y}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} \mathbb{V}_{f | \mathcal{D}_t \cup \mathcal{D}^*} \kappa_{u,f}(\boldsymbol{\theta}, \boldsymbol{\theta}') \\ &= 2T \left( \frac{\tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')} , \sqrt{\frac{\sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)}{\sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') + \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)}} \right), \end{aligned} \quad (21)$$

respectively. Above  $T(\cdot, \cdot)$  denotes the Owen’s  $T$  function (Owen, 1956, 1980) and

$$\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) + \tau_t^2(\boldsymbol{\theta}'; \boldsymbol{\theta}^*) - 2\omega_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*), \quad (22)$$

$$\tau_t^2(\boldsymbol{\theta}_\bullet; \boldsymbol{\theta}^*) = \omega_t(\boldsymbol{\theta}_\bullet, \boldsymbol{\theta}_\bullet; \boldsymbol{\theta}^*), \quad (23)$$

$$\omega_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = c_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) [c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \boldsymbol{\Lambda}^*]^{-1} c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}'). \quad (24)$$

Note that above  $\mathcal{E}_{t+b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  and  $E_{t+b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  both depend on  $\mathbf{y}^*$  and  $\boldsymbol{\theta}^*$  via  $\mathcal{D}^*$  although this is not explicitly shown in the notation.

Another strategy would be to select such  $\boldsymbol{\theta}^*$  that minimises the future variance of the log-MH ratio  $\mathbb{V}_{f | \mathcal{D}_t \cup \mathcal{D}^*} \log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_{t+b}^2(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  as this acquisition function conveniently does not depend on the unknown  $\mathbf{y}^*$ . Interestingly,  $\mathbb{V}_{f | \mathcal{D}_t \cup \mathcal{D}^*} \log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')$  and all three acquisition functions of Proposition 2 share the same global minimiser:

**Proposition 3** *The global minimum  $\boldsymbol{\theta}_{\text{opt}} \in \Theta^b$  for the variance of the log-MH ratio  $\mathbb{V}_{f|\mathcal{D}_t \cup \mathcal{D}^*} \log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')$ , for the expected conditional error  $L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  with any  $u \in [0, 1]$ , for the expected unconditional error  $L_t^E(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  and also for the expected variance  $L_t^{\text{v},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  of  $\kappa_{u,f}$  with any  $u \in [0, 1]$ , is given by*

$$\boldsymbol{\theta}_{\text{opt}} \in \arg \max_{\boldsymbol{\theta}^* \in \Theta^b} \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*). \quad (25)$$

We interpret  $\arg \max_{\boldsymbol{\theta}^* \in \Theta^b} \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  as a set because the minimiser in (25) may not be unique. When  $b = 1$  and  $\Theta = \prod_{i=1}^p [a_i, b_i]$  where we allow  $a_i = -\infty$  and  $b_i = \infty$ , the optimisation task (25) could be simplified by replacing  $\Theta$  with some bounded set  $\tilde{\Theta} \subset \Theta$  located around  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$ . For example, we can choose

$$\tilde{\Theta} = \prod_{i=1}^p [\max\{\min\{\theta_i, \theta'_i\} - cl_i, a_i\}, \min\{\max\{\theta_i, \theta'_i\} + cl_i, b_i\}], \quad (26)$$

where  $l_i$ 's denote the lengthscales of the GP covariance function and  $c > 0$  controls the size of the set. Of course, the set (26) is not guaranteed to contain the global optimum unless  $c$  is large enough. Another tempting choice is  $\tilde{\Theta} = \{\boldsymbol{\theta}, \boldsymbol{\theta}'\}$  which simplifies the implementation and facilitates fast computations because no auxiliary global optimisation is then needed.

## 4. Interpretations of GP-MH

In Section 4.1 we reason that the standard and noisy/pseudo-marginal MH samplers are special cases of the GP-MH framework when an uninformative GP prior is used. In Section 4.2 we show a concrete relation between GP-MH and the BLFI frameworks. This motivates an alternative, two-stage implementation of GP-MH called MH-BLFI. In Section 4.3 we argue that when the randomness of the MH sampler in Algorithm 1 is considered fixed, the sampler can be treated as a deterministic mapping whose input is the likelihood function and whose output is a (multi)set of parameter values. The uncertainty in the output due to the use of the GP model in the place of the exact likelihood function can then be probabilistically quantified in principle. GP-MH can be viewed as a heuristic yet tractable estimate for the MH output under this conceptual ‘‘Bayesian’’ MH setting.

### 4.1 Standard and noisy/pseudo-marginal MH as special cases of GP-MH

Throughout this section the log-likelihood function  $f$  is modelled with a white noise process. That is, the GP prior has mean  $m_0(\boldsymbol{\theta}) = 0$  and covariance function  $k_\phi(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_s^2 \mathbb{1}_{\boldsymbol{\theta}=\boldsymbol{\theta}'}$  with  $\phi = \sigma_s^2 > 0$ . We further assume  $\sigma_s^2$  is set arbitrarily large.

First, suppose  $\sigma_n^2(\boldsymbol{\theta}) = 0$  for all  $\boldsymbol{\theta}$  and  $0 \leq \varepsilon < 1/2$ . The evaluation at the current point  $\boldsymbol{\theta}$  is automatically reused from the previous iteration via the GP model and is not recomputed. If the proposed point  $\boldsymbol{\theta}'$  has not been visited before, GP-MH algorithm needs to evaluate at  $\boldsymbol{\theta}'$  because otherwise the (un)conditional error would remain  $1/2$ . All of our acquisition functions will select  $\boldsymbol{\theta}'$  in this case and the (un)conditional error becomes 0 after this exact evaluation, see Section C.1.2 for further details on this. If  $\boldsymbol{\theta}'$  has been visited before, its evaluation is also automatically reused via the GP and the (un)conditional error is 0. Hence, we conclude that GP-MH reverts to standard MH where the GP model has no role but to cache the old evaluations for reuse.

Suppose next that the evaluations are noisy so that  $\sigma_n^2(\boldsymbol{\theta}) > 0$  for all  $\boldsymbol{\theta}$  and  $0 < \varepsilon < 1/2$ . Noisy/pseudo-marginal MH is then technically not a special case of Algorithm 3 but falls under a slightly generalised version of our framework where one (and just one) noisy evaluation at  $\boldsymbol{\theta}'$  is acquired even when the resulting (un)conditional error is not smaller than  $\varepsilon$ . The GP model again caches the old evaluations. A disagreement however arises if the proposal  $\boldsymbol{\theta}'$  has been visited before. Then noisy/pseudo-marginal MH would use a new noisy evaluation at  $\boldsymbol{\theta}'$  whereas the old evaluation at  $\boldsymbol{\theta}'$  would also be taken into account via the GP model in GP-MH. This disagreement however occurs with probability zero in the common case where  $\Theta$  is uncountable and the proposal  $q$  does not contain any point masses. Finally, we observe that Algorithm 3 with the white noise prior and  $0 < \varepsilon < 1/2$  can be seen as a variant of the algorithm by Korattikara et al. (2014) where Bayesian analysis is used instead of their frequentist hypothesis testing approach to determine each MH accept/reject decision.

## 4.2 The relation between GP-MH and BLFI

Suppose no new log-likelihood evaluations are needed in Algorithm 3 after some iteration  $i$  because any possible later transition can be done confidently enough based on the current GP model. From iteration  $i$  onwards Algorithm 3 then acts as an exact MH sampler that targets a density proportional to

$$\text{med}_{f|\mathcal{D}_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_i(\boldsymbol{\theta})}. \quad (27)$$

This fact follows because in this case  $\mathcal{D}_t$  does not change after iteration  $i$  and because

$$\hat{\gamma} = \text{med}_{f|\mathcal{D}_t}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')) = e^{\mu_i(\boldsymbol{\theta}, \boldsymbol{\theta}')} = \frac{\pi(\boldsymbol{\theta}')e^{m_i(\boldsymbol{\theta}')}q(\boldsymbol{\theta}|\boldsymbol{\theta}')}{\pi(\boldsymbol{\theta})e^{m_i(\boldsymbol{\theta})}q(\boldsymbol{\theta}'|\boldsymbol{\theta})} = \frac{\text{med}_{f|\mathcal{D}_t}(\tilde{\pi}_f(\boldsymbol{\theta}'))q(\boldsymbol{\theta}|\boldsymbol{\theta}')}{\text{med}_{f|\mathcal{D}_t}(\tilde{\pi}_f(\boldsymbol{\theta}))q(\boldsymbol{\theta}'|\boldsymbol{\theta})}.$$

This situation occurs also if a hard threshold  $t_{\max}$  on the maximum number of log-likelihood evaluations, after which no more evaluations are collected even if some later proposed transition would be unconfident, is introduced to Algorithm 3.

### 4.2.1 GP-MH AS A BLFI METHOD WITH A STOCHASTIC ACQUISITION FUNCTION

The above remark suggests an alternative, two-stage variant of GP-MH: 1) Run Algorithm 3 until  $t_{\max}$  evaluations have been gathered or some other stopping criterion is met. 2) Treat the estimate (27) based on the evaluations collected in stage 1 as an approximation to the unnormalised posterior and plug it in to some standard MCMC sampler. We write this idea in the form of Algorithm 4 which in fact is a special case of BLFI (compare Algorithm 4 to Algorithm 2) and we call it MH-BLFI. The tail-recursive procedure NEXTUNCONFIDENT-TRANSITION together with line (10) is interpreted as a sequential stochastic strategy for selecting the evaluation locations in the BLFI framework. The recursion never completes if the condition  $E_{t,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \leq \varepsilon$  holds for all  $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$  so in practice Algorithm 4 would need to be further modified to terminate prematurely if the recursion becomes too deep.

A key difference between GP-MH in Algorithm 3 and MH-BLFI in Algorithm 4 is that in the former approach the GP model is continuously refined. A potential concern of the former approach is that its convergence as an approximate MH sampler is problematic to

---

**Algorithm 4** GP-MH reformulated in BLFI framework (MH-BLFI)

---

**Input:** Prior  $\pi(\boldsymbol{\theta})$ , GP model for  $f$ , error tolerance  $\varepsilon$ , no. initial evaluations  $t_{\text{init}}$ , no. total evaluations  $t_{\text{max}}$ , initial point  $\boldsymbol{\theta}^{(0)}$ , proposal  $q(\boldsymbol{\theta}' | \boldsymbol{\theta})$ , no. MCMC samples  $s_{\text{MCMC}}$

**Output:** Samples  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(s_{\text{MCMC}})}$

1-6: Obtain initial GP and  $\mathcal{D}_{t_{\text{init}}}$  ▷ Lines 1-6 are the same as those in Algorithm 3.

7: Set  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}^{(0)}$  and sample  $\boldsymbol{\theta}' \sim q(\cdot | \boldsymbol{\theta})$  and  $u \sim \mathcal{U}([0, 1])$

8: **for**  $t = t_{\text{init}} + 1 : t_{\text{max}}$  **do**

9:     Set  $(\boldsymbol{\theta}, \boldsymbol{\theta}', u) \leftarrow \text{NEXTUNCONFIDENTTRANSITION}(\boldsymbol{\theta}, \boldsymbol{\theta}', u)$

10:    Obtain  $\boldsymbol{\theta}^*$  as a solution to (25) using  $(\boldsymbol{\theta}, \boldsymbol{\theta}', u)$  ▷ As in Algorithm 3.

11:    Set  $y^* \leftarrow$  log-likelihood evaluation at  $\boldsymbol{\theta}^*$

12:    Set  $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup \{(y^*, \boldsymbol{\theta}^*)\}$

13:    Refit GP using  $\mathcal{D}_t$

14: **end for**

15: Sample  $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(s_{\text{MCMC}})}$  from (27) with MCMC ▷ Alternatively, use (29).

16: **procedure**  $\text{NEXTUNCONFIDENTTRANSITION}(\boldsymbol{\theta}, \boldsymbol{\theta}', u)$

17:    **if**  $E_{t, \hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') > \varepsilon$  **then**

18:      **Return**  $(\boldsymbol{\theta}, \boldsymbol{\theta}', u)$

19:    **else if**  $\hat{\gamma}(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq u$  **then** ▷ Proposal  $\boldsymbol{\theta}'$  is accepted.  $\hat{\gamma}$  computed using (15).

20:       $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}'$

21:    **end if** ▷ Only the latest accepted  $\boldsymbol{\theta}$  needs to be stored.

22:    Sample  $\boldsymbol{\theta}' \sim q(\cdot | \boldsymbol{\theta})$  and  $u \sim \mathcal{U}([0, 1])$

23:    **Return**  $\text{NEXTUNCONFIDENTTRANSITION}(\boldsymbol{\theta}, \boldsymbol{\theta}', u)$

24: **end procedure**

---

assess because the target density is slightly changing. On the other hand, the fact that the GP surrogate can no longer be improved in the second stage of the latter approach can also be problematic. See also Llorente et al. (2021, Section 3) for some generic discussion on such approaches. Another important difference is that Algorithm 3 is run for  $i_{\text{MH}}$  iterations and the number of log-likelihood evaluations is not known in advance whereas Algorithm 4 uses  $t_{\text{max}}$  evaluations (or possibly less if an additional stopping criterion is used).

#### 4.2.2 ROBUST ESTIMATOR OF THE UNNORMALISED POSTERIOR

An explicit estimator for the unnormalised posterior  $\tilde{\pi}_f(\boldsymbol{\theta})$  is needed in a two-stage approach such as MH-BLFI in Algorithm 4. The choice of this estimator can be based on Bayesian decision theory. As discussed by Järvenpää et al. (2020), the (marginal) median  $\hat{p}_1(\boldsymbol{\theta}) = \text{med}_{f | \mathcal{D}_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_t(\boldsymbol{\theta})}$ , which was also already shown as (27), minimises the expected  $L^1$ -loss so that  $\hat{p}_1(\boldsymbol{\theta}) = \arg \min_{\tilde{p}} \mathbb{E}_{f | \mathcal{D}_t} \tilde{l}_1(\tilde{\pi}_f, \tilde{p})$  where  $\tilde{l}_1(\tilde{\pi}_f, \tilde{p}) = \int_{\Theta} |\tilde{\pi}_f(\boldsymbol{\theta}) - \tilde{p}(\boldsymbol{\theta})| d\boldsymbol{\theta}$ . This estimator is intuitive and easy to compute.

Similarly to standard MCMC methods, the parameter regions far in the posterior tails are typically not extensively explored when Algorithm 3 and 4 are used and so  $\mathcal{D}_t$  mostly contains points in the highest density region. Consequently, the uncertainty of the likelihood function can remain large in such unexplored tail regions and the resulting GP-based

estimator for  $\tilde{\pi}_f(\boldsymbol{\theta})$ , such as the median (27), can unintuitively have a non-negligible value there. This can produce difficult target densities for MCMC in the second stage of a two-stage approach such as MH-BLFI. The issue has also been observed by Fielding et al. (2011); Drovandi et al. (2018) and is illustrated in our set-up in Figure 2: The estimated SL posterior, the red line in Figure 2b, has two distant modes. The possibility of the second mode at  $\theta = 30$  cannot be excluded based on the GP model fitted to the 9 noisy evaluations.




Figure 2: (a) Log-SL of a simple 1D toy problem is modelled using a GP with zero-mean prior. Red line (red shaded region) shows the GP posterior mean (95% credible interval), black dots/lines the log-SL evaluations with corresponding observation errors. (b) The (rescaled) SL, which equals the (unnormalised) SL posterior when a uniform prior on  $[-20, 30]$  is used, then follows a log-GP. Red line shows the marginal median estimate (27) and dashed blue line the marginal mode (29). The lack of evaluations in  $\theta \in [10, 30]$  causes large uncertainty in this region as shown by the shaded red region.

In the practical application of standard MCMC methods all unvisited tail regions are implicitly and inevitably neglected as unlikely as only the visited locations may belong to the final set of samples. MH-BLFI however requires different, more explicit solution. Ideally, the above practical issue is alleviated in MH-BLFI by incorporating more prior information regarding the possibility of multimodality to the GP model or by meaningfully altering the original prior density  $\pi(\boldsymbol{\theta})$  but this is cumbersome. Here we instead propose an estimator for the unnormalised posterior that is shrunk towards zero in regions with large uncertainty.

Consider the loss function

$$\tilde{l}_g(\tilde{\pi}_f, \tilde{p}) \triangleq \int_{\Theta} (1 - g(\boldsymbol{\theta}))(\tilde{\pi}_f(\boldsymbol{\theta}) - \tilde{p}(\boldsymbol{\theta}))\mathbb{1}_{\tilde{\pi}_f(\boldsymbol{\theta}) \geq \tilde{p}(\boldsymbol{\theta})} + g(\boldsymbol{\theta})(\tilde{p}(\boldsymbol{\theta}) - \tilde{\pi}_f(\boldsymbol{\theta}))\mathbb{1}_{\tilde{p}(\boldsymbol{\theta}) > \tilde{\pi}_f(\boldsymbol{\theta})} d\boldsymbol{\theta}, \quad (28)$$

where  $g : \Theta \rightarrow (0, 1)$  is a weight function. Clearly, the choice  $g(\boldsymbol{\theta}) = 1/2$  gives the  $L^1$  loss. By changing the order of expectation and integration and then using Proposition 2.5.5 in Robert (2007), we can see that the posterior expected loss  $\mathbb{E}_{f|\mathcal{D}_t} \tilde{l}_g(\tilde{\pi}_f, \tilde{p})$  is minimised when  $\tilde{p}(\boldsymbol{\theta})$  is the  $(1 - g(\boldsymbol{\theta}))$ -percentile of the log-Normal distribution of  $\tilde{\pi}_f(\boldsymbol{\theta}) | \mathcal{D}_t$  for (almost) all  $\boldsymbol{\theta} \in \Theta$ . The resulting estimator is hence  $\pi(\boldsymbol{\theta}) \exp(m_t(\boldsymbol{\theta}) + \Phi^{-1}(1 - g(\boldsymbol{\theta}))s_t(\boldsymbol{\theta}))$ . If we exceptionally allow  $g$  to depend on the posterior of  $f$  and choose  $g(\boldsymbol{\theta}) = \Phi(s_t(\boldsymbol{\theta}))$  so that the loss function (28) penalises large posterior estimates in the regions with large uncertainty, we obtain the estimator

$$\hat{p}_g(\boldsymbol{\theta}) = \text{mode}_{f|\mathcal{D}_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_t(\boldsymbol{\theta}) - s_t^2(\boldsymbol{\theta})}, \quad (29)$$



which is the marginal mode. Estimator (29) behaves similarly as (27) in the highest density region where typically  $s_t^2(\boldsymbol{\theta}) \approx 0$  or  $s_t^2(\boldsymbol{\theta}) \ll |m_t(\boldsymbol{\theta})|$  but meaningfully shrinks its value towards 0 in regions with large uncertainty as seen in Figure 2b (the dashed blue line)<sup>1</sup>.

### 4.3 “Bayesian” MH sampler

We take here a more conceptual approach than in other sections and again revisit the MH sampler in Algorithm 1. Samples from  $q(\boldsymbol{\theta}' | \boldsymbol{\theta})$  are often obtained using a relation  $\boldsymbol{\theta}' = g(\boldsymbol{\theta}, \mathbf{r})$ , where  $g : \Theta \times \mathbb{R}^r \rightarrow \Theta$  is a known function and  $\mathbf{r}$  follows some standard distribution. For example, the independent MH sampler is obtained using  $g(\boldsymbol{\theta}, \mathbf{r}) = \mathbf{r}$  and the Gaussian proposal  $q(\boldsymbol{\theta}' | \boldsymbol{\theta}) = \mathcal{N}_p(\boldsymbol{\theta}' | \boldsymbol{\theta}, \boldsymbol{\Sigma})$  follows as  $g(\boldsymbol{\theta}, \mathbf{r}) = \boldsymbol{\theta} + \mathbf{r}, \mathbf{r} \sim \mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$ . Although we could proceed more generally, in the following we assume the relation  $\boldsymbol{\theta}' = g(\boldsymbol{\theta}, \mathbf{r}) = \boldsymbol{\theta} + \mathbf{r}$ , where  $\mathbf{r}$  follows some absolute continuous density (e.g. Gaussian with a non-singular covariance matrix  $\boldsymbol{\Sigma}$ ). We disregard here any intricate issues related to the convergence and initialisation of MH. That is, in order to gain an alternative theoretical interpretation of the GP model in the context of an MH sampler, we consider an ideal scenario where the initial point  $\boldsymbol{\theta}^{(0)}$  is located in the highest density region, a suitable proposal  $q$  is immediately available and a single chain (length  $n$ , no burn-in) long enough to produce a negligible Monte Carlo error is run.

A preliminary key observation is that instead of drawing  $u^{(i)}$  and  $\boldsymbol{\theta}'^{(i)}$  at each iteration  $i$  of Algorithm 1,  $u^{(i)}$  and  $\mathbf{r}_i$  for  $i = 1, \dots, n$  can be pre-generated. From now on we hence suppose  $u^{(i)}$ 's and  $\mathbf{r}_i$ 's are fixed so we can exceptionally treat MH as a deterministic algorithm (or mapping) whose input is the likelihood function and output a (multi)set of  $n$  samples. The GP posterior  $f | \mathcal{D}_t \sim \mathcal{GP}(m_t(\boldsymbol{\theta}), c_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$  of the log-likelihood then induces a probability distribution over the output, the  $n$  samples. The corresponding estimate  $\hat{h}_n$  in (3) similarly follows a probability distribution induced by the GP. We call this unusual approach “Bayesian” MH. In this approach the posterior uncertainty of  $f$  is accounted probabilistically but the sampling error due to the finite sample size  $n$  is not. This makes Bayesian MH fundamentally different from the related Bayesian quadrature methods (O’Hagan, 1991; Rasmussen and Ghahramani, 2003; Briol et al., 2019).

The proposal is either accepted or the current point is kept at each iteration of MH. Hence, the possible states of MH sampler at iteration  $i$  are  $\mathcal{S}_i \triangleq \{\boldsymbol{\theta}^{(0)} + \sum_{j=1}^i e_j \mathbf{r}_j \mid e_j \in \{0, 1\} \forall j \in \{1, \dots, i\}\}$  so that  $\mathcal{S}_0 = \{\boldsymbol{\theta}^{(0)}\}$ ,  $\mathcal{S}_1 = \{\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_1\}$ ,  $\mathcal{S}_2 = \{\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_1, \boldsymbol{\theta}^{(0)} + \mathbf{r}_2, \boldsymbol{\theta}^{(0)} + \mathbf{r}_1 + \mathbf{r}_2\}$  and so on. We also have  $\mathcal{S}_i \subset \mathcal{S}_{i+1}, i \geq 0$ . From now on we assume  $|\mathcal{S}_i| = 2^i$  for all  $0 \leq i \leq n$  which would happen with probability 1 anyway as we already assumed that  $\mathbf{r}_i$ 's are realisations of an absolute continuous density. We denote the (random) state of Bayesian MH at iteration  $i$  as  $\boldsymbol{\theta}_{(i)} \in \mathcal{S}_i$ . The Bayesian MH forms itself a discrete-time process with finite state space  $\mathcal{S}_i$  whose cardinality grows exponentially as a function of iteration  $i$ . The process is not Markovian in general so that the probability mass function of each realisation only satisfies  $p(\boldsymbol{\theta}_{(0)}, \dots, \boldsymbol{\theta}_{(n)} | \mathcal{D}_t) = \prod_{i=1}^n p(\boldsymbol{\theta}_{(i)} | \boldsymbol{\theta}_{(i-1)}, \dots, \boldsymbol{\theta}_{(0)}, \mathcal{D}_t) p(\boldsymbol{\theta}_{(0)})$ , where the initial probability mass function is  $p(\boldsymbol{\theta}_{(0)}) = \mathbb{1}_{\boldsymbol{\theta}_{(0)} = \boldsymbol{\theta}^{(0)}}$ . For example, the posterior probability that the true MH chain would first stay at the initial point  $\boldsymbol{\theta}^{(0)}$  and then move to the proposed point would be  $p(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_2 | \mathcal{D}_t) = \mathbb{P}_{f | \mathcal{D}_t}(\gamma_f(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_1) <$

1. In Appendix D we show that a more suitable GP model or an additional evaluation near the right boundary also remove the problematic, second mode in this particular 1D case.

$u^{(1)}, \gamma_f(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_2) \geq u^{(2)}$ ) which could be computed using the bivariate Gaussian CDF. There are exactly  $2^n$  paths the true chain can take and the probability of other paths is hence 0. For example,  $p(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(0)} + \mathbf{r}_1) = 0$  since a transition from  $\boldsymbol{\theta}^{(0)}$  to  $\boldsymbol{\theta}^{(0)} + \mathbf{r}_1$  can only happen at iteration  $i = 1$ .

The expectation of  $\hat{h}_{n+1} = \sum_{i=0}^n h(\boldsymbol{\theta}_{(i)}) / (n + 1)$  with respect to  $f | \mathcal{D}_t$  is computed as

$$\begin{aligned} \mathbb{E}_{f | \mathcal{D}_t}(\hat{h}_{n+1}) &= \sum_{(\boldsymbol{\theta}_{(0)}, \dots, \boldsymbol{\theta}_{(n)}) \in \prod_{i=0}^n \mathcal{S}_i} \frac{1}{n+1} \sum_{i=0}^n h(\boldsymbol{\theta}_{(i)}) p(\boldsymbol{\theta}_{(0)}, \dots, \boldsymbol{\theta}_{(n)} | \mathcal{D}_t) \\ &= \frac{1}{n+1} \sum_{i=0}^n \sum_{\boldsymbol{\theta}_{(i)} \in \mathcal{S}_i} h(\boldsymbol{\theta}_{(i)}) p(\boldsymbol{\theta}_{(i)} | \mathcal{D}_t). \end{aligned}$$

A formula for the variance of  $\hat{h}_{n+1}$  can also be derived. Unfortunately, these computations seem to require repeated evaluations of multivariate Gaussian CDF and would not scale better than  $\mathcal{O}(2^n)$  in any case. Even if some realisations could be neglected as impossible (e.g. chains that lead outside the support of the prior density) or extremely unlikely based on the GP, the computations remain intractable. Generating a sample path of  $(\boldsymbol{\theta}_{(0)}, \dots, \boldsymbol{\theta}_{(n)})$  is more feasible as this requires only an iterative generation of a GP sample path. The resulting  $\mathcal{O}(n^3)$  cost still limits this approach to short chains only. GP approximations (see e.g. Wilson et al., 2020) could be used to bypass the cubic cost but we do not consider them here.

We may now view Algorithm 3, as well as the related GPS-ABC algorithm by Meeds and Welling (2014), as a heuristic yet tractable approach for constructing an estimator for the true MH chain. In these algorithms the most probable decision, either acceptance or rejection of the proposed point, is selected (un)conditionally on  $u^{(i)}$  at each iteration  $i$  without acknowledging possible future accept/reject decisions or the locations visited during iterations  $0, 1, \dots, i - 2$ . Perhaps a more natural estimator would be the most probable set of samples but it also appears intractable. The uncertainty of the resulting set of samples (or of the estimate  $\hat{h}_n$ ) due to the GP posterior is not explicitly quantified.

An ideal experimental design strategy for Bayesian MH would be such that minimises the expected uncertainty of the samples or of  $\hat{h}_n$ . Such a strategy is also intractable even under the one-step ahead simplification. Furthermore, this strategy would presumably feature similar practical challenges as mentioned in Section 1, e.g. the need for accurate global modelling. The tractable sequential strategies developed in Section 3 instead minimise the expected uncertainty regarding the current accept/reject decision without explicitly accounting for the aforementioned final goal of the inference.

#### 4.4 Summary of theoretical analysis

We briefly outline the two most important theoretical results regarding our GP-MH algorithm. The first key result is that the optimal evaluation location  $\boldsymbol{\theta}_{\text{opt}}$  of (25) does not generally coincide with  $\boldsymbol{\theta}$  or  $\boldsymbol{\theta}'$  unlike one might first intuitively expect, see Appendix C.1 for our detailed analysis on this. This motivates the experimental design strategies of Section 3 and suggests that restricting the optimisation in (25) to  $\tilde{\Theta} = \{\boldsymbol{\theta}, \boldsymbol{\theta}'\}$  may not produce the best possible sample-efficiency. We further investigate this empirically in Section 5.

In Section C.2 we analyse the number of log-likelihood evaluations needed at an individual iteration of GP-MH. The results there demonstrate a potential shortcoming of GP-MH: Some individual steps of the algorithm may need even hundreds of evaluations. However, these results concern the worst case situation and are not representative of practice where most accept/reject decisions are done based solely on the GP model and already during the early iterations. The GP-MH algorithm may still occasionally get “stuck” (somewhat similarly as pseudo-marginal MH but for a completely different reason) especially if  $\varepsilon$  is chosen “too small”. The numerical experiments in Section 5 show that sample-efficiency comparable to BLFI is reached with suitable choices of  $\varepsilon$ .

## 5. Numerical experiments

In this section we investigate the effect of the tolerance parameter  $\varepsilon$  and the developed sequential experimental design strategies on the quality of the resulting posterior approximation. We compare our GP-MH and MH-BLFI implementations and also consider a BLFI implementation with the theoretically well-motivated and best-performing integrated median interquantile range (IMIQR) strategy by Järvenpää et al. (2021). We consider three scenarios: 1) synthetically constructed log-densities corrupted with additive Gaussian noise (Section 5.1), 2) SL inference for simulator-based models (Section 5.2), 3) likelihood-free generalised Bayesian inference (Section 5.3). This allows us to understand how the algorithms perform both in ideal circumstances and in more realistic situations where the GP modelling assumptions are to some extent violated.

Our GP-MH and MH-BLFI implementations both use the same GP surrogate and mainly differ in how the final posterior approximation is formed, as discussed in Section 4.2. Instead of using MH-BLFI with pre-determined number of evaluations as in Algorithm 4, it is implemented here so that a separate MH sampler targeting the robust mode-based estimate (29) is run at various stages during Algorithm 3. This approach simplified our experiments and reduced their total computational cost. We use the unconditional error (18) and the two strategies of Section 3 for selecting the evaluations: 1) “EPoE” which stands for expected probability of error and requires solving (25) over the set (26) with  $c = 3/4$ , and 2) “EPoEr” where the extra “r” informs that the optimisation in (25) is restricted to  $\tilde{\Theta} = \{\boldsymbol{\theta}, \boldsymbol{\theta}'\}$ . We also consider a baseline “naive”, where the new evaluation location  $\boldsymbol{\theta}^*$  is the current point  $\boldsymbol{\theta}$  with probability 0.5 and the proposed point  $\boldsymbol{\theta}'$  otherwise.

In each experiment, the initial point  $\boldsymbol{\theta}^{(0)}$  is chosen to be near (but outside) the highest density region to represent the scenario where some weak information about the location is available, e.g. as a result of pilot runs. We use a Gaussian proposal whose covariance matrix is updated adaptively as mentioned in Section 2.7. The first quarter of the (approximate) MCMC samples is always neglected as burn-in. We use the same GP model of Section 2.5 for all of our experiments. In particular, we use basis functions  $1, \theta_j, \theta_j^2$  for each dimension  $j$  and we set  $\mathbf{b} = \mathbf{0}$  and  $B_{jk} = 30^2 \mathbf{1}_{j=k}$  although this is likely unideal. We assume that the log-likelihood is a smooth function on the highest density region and use the squared exponential covariance function  $k_\phi(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_s^2 \exp(-\sum_{j=1}^p (\theta_j - \theta'_j)^2 / (2l_j^2))$ . We further use relatively uninformative hyperpriors for the GP hyperparameters  $\boldsymbol{\phi} = (\sigma_s^2, l_1, \dots, l_p)$  (and  $\sigma_n^2$  in Section 5.1 and 5.3). The hyperparameters are re-estimated immediately after each new log-likelihood evaluation when  $t \leq 300$  and after every 10th evaluation otherwise.

We summarise the posterior approximation accuracy mainly via the average total variation distance over the coordinate-wise marginal densities as compared to the ground-truth. In the 2D cell biology experiment we instead use the joint total variation. That is, in the former case we define  $\text{TV}(\pi, \pi') \triangleq \sum_{j=1}^p \int_{\Theta_j} |\pi(\theta_j) - \pi'(\theta_j)| d\theta_j / (2p)$  and in the latter case  $\text{TV}(\pi, \pi') \triangleq \int_{\Theta} |\pi(\boldsymbol{\theta}) - \pi'(\boldsymbol{\theta})| d\boldsymbol{\theta} / 2$  where  $\pi, \pi'$  are pdfs both defined over  $\Theta = \prod_{j=1}^p \Theta_j \subset \mathbb{R}^p$  and  $\pi(\theta_j), \pi'(\theta_j)$  denote their marginals. We chose these criteria to be in line with previous work and because their values are easily approximated using the MCMC output, are not overly sensitive to inevitable small errors especially in the tails and are easy to interpret (as the overlap between the densities in the scale  $[0, 1]$ ). Each run of the algorithms is repeated 100 times (50 times for cell biology experiment) with different realisation of randomness to assess the variability. The experiments were performed using MATLAB 2022a. Some GP functionality was taken from GPstuff 4.7 (Vanhatalo et al., 2013).

### 5.1 Synthetic log-densities

We first consider three 6D densities from Järvenpää et al. (2021) with different characteristics: a Gaussian density called ‘Simple’, a banana-shaped density ‘Banana’ and a multimodal density ‘Multimodal’. The variance of the log-density evaluations  $\sigma_n^2(\boldsymbol{\theta})$  is here constant and estimated together with the GP hyperparameters  $\boldsymbol{\phi}$ . We use  $\sigma_n = 1$  for Banana and Multimodal and  $\sigma_n = 2$  for the Simple log-density. Further details of the target densities are given in Appendix F and some additional results in Appendix G.1. The number of initial evaluations is  $t_{\text{init}} = 10$ . Algorithm 3 is run for  $i_{\text{MH}} = 10^5$  iterations.

Figure 3 show the median accuracy and the variability of the final posterior approximation. A hypothetical best possible algorithm that always produces an exact posterior ( $\text{TV} = 0$ ) without any log-density evaluations would appear in the left lower corner of the figure. We see that EPoE produces the best sample-efficiency, the naive method is the worst and EPoEr is roughly halfway between them. Decreasing  $\varepsilon$  leads to more accurate posterior approximations as expected. Interestingly, a fairly large tolerance ( $\varepsilon \approx 0.3$ ) is already sufficient to produce reasonable approximations. In this case EPoE needs at most a few hundred log-density evaluations. Banana log-density is more challenging than the other two as all methods face some challenges in estimating its long tails. GP-MH and MH-BLFI implementations produce similar results which is because most log-density evaluations are already collected during the “burn-in”, as seen in Figure G.2 of Appendix. The results also show that EPoE produces similar or slightly worse accuracy as BLFI (IMIQR) with large values of  $\varepsilon$  and performs slightly better with low  $\varepsilon$ .

Typical examples of selected evaluation locations are shown in Figure 4. EPoE tends to select slightly more diverse evaluation locations as the two other methods. On the other hand, the evaluation locations of IMIQR are even more diverse and, despite the fairly large number of initial locations (20 for Simple and 50 for Banana), some of its evaluations occur at the boundary of the parameter space. When the initial point  $\boldsymbol{\theta}^{(0)}$  is far from the highest density region, evaluations intuitively occur on a path connecting the initial point and the highest density region as most clearly seen in the top left case of Figure 4.

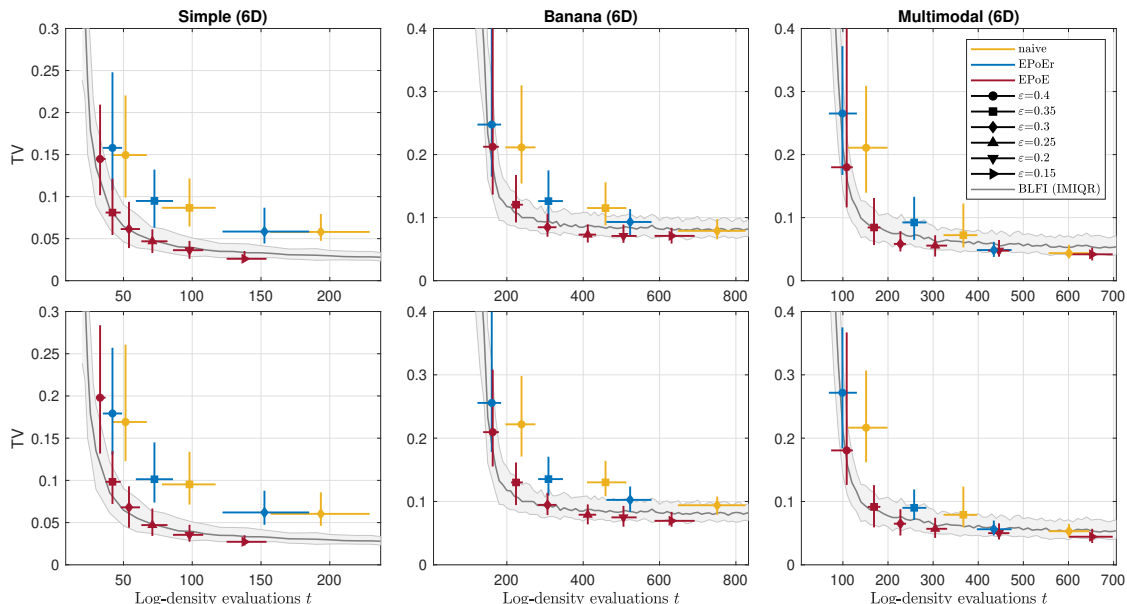


Figure 3: Accuracy of the marginal posterior approximation as a function of the used log-density evaluations at the final iteration  $i_{\text{MH}} = 10^5$ . The horizontal and vertical lines show the middle 75% computed over 100 repeated runs and the marker in the middle shows the corresponding (marginal) median. The grey lines/area show the median and the middle 75% for BLFI with IMIQR design strategy over 100 repeats. *Top row* shows GP-MH and the *bottom row* the corresponding results by MH-BLFI.

## 5.2 SL inference

We next consider simulator-based models where SL is used as the target likelihood, see Section A.1 of Appendix for details on SL. Although our methodology is particularly useful for expensive models, we here nevertheless use models that are not highly costly (although not very cheap either). This allows us to compare our results directly to reasonable ground-truth posteriors obtained via extensive computation using SL-MCMC (Price et al., 2018). For simplicity, we always evaluate log-SL using fixed number of repeated simulations  $N$ .

We use the bootstrap to estimate the noise variance  $\sigma_n(\boldsymbol{\theta})$  at the evaluated locations in  $\mathcal{D}_t$  but this does not work for EPoE which needs estimates of  $\sigma_n(\boldsymbol{\theta})$  at each  $\boldsymbol{\theta}$ . We first approximated  $\sigma_n(\boldsymbol{\theta})$  for EPoE with a constant obtained near the MAP parameter value but this resulted underestimated  $\sigma_n(\boldsymbol{\theta}^*)$  and hence overestimated  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  in the tails which further caused overexploration of the tails. As a heuristic solution we set  $\sigma_n(\boldsymbol{\theta}) = 0.1$  as if the evaluations were almost noiseless. More realistic estimates could possibly be obtained by modelling  $\sigma_n(\boldsymbol{\theta})$  as a function of  $\boldsymbol{\theta}$  (nearest neighbour interpolation was used by Acerbi (2020) for a similar goal but this approach would make  $\sigma_n(\boldsymbol{\theta})$ —and consequently EPoE—discontinuous) but our simple approach already produced improvements over the naive approach. EPoEr was observed to be insensitive for this choice. We leave more detailed analysis on modelling the noise variance, and the log-likelihood function itself, an important topic for future work.

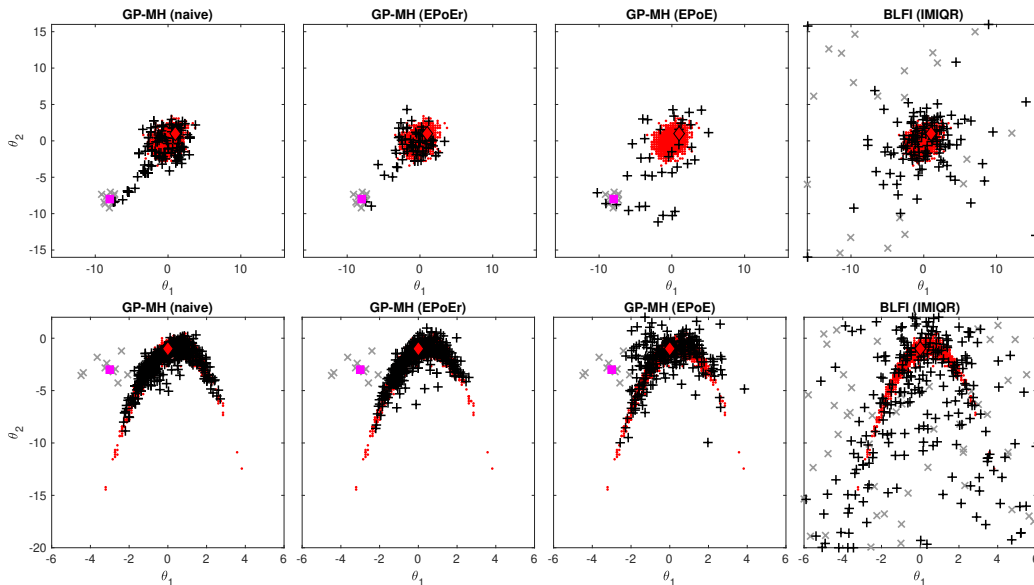


Figure 4: Typical realisations of the log-density evaluation locations in the Simple (*top row*) and Banana (*bottom row*) experiments. The evaluation locations are projected to the first two components. Red dots in the background illustrate the true target density, the grey crosses ( $\times$ ) the  $t_{\text{init}}$  initial evaluation locations and the black plus signs ( $+$ ) the evaluation locations. The magenta square shows the initial location.

### 5.2.1 THETA-RICKER MODEL

We consider theta-Ricker model, see e.g. Polansky et al. (2009) and references therein for background. In this model  $N_t$  denotes the number of individuals in a population (or population density) at time  $t$  which evolves according to the discrete time process

$$N_{t+1} = rN_t \exp(-\log(r)(N_t/K)^\theta + \varepsilon_t),$$

for  $t = 1, \dots, T$ . Parameter  $K$  indicates the population size when the growth rate  $r$  goes to zero and  $\theta$  further controls the form of the growth rate. The Ricker model, a common LFI test problem that we also consider in Appendix G.2, is a special case with  $\theta = 1$  and  $K = \log(r)$ . Gaussian process noise model  $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$  is assumed. A noisy measurement  $x_t$  of the population size  $N_t$  is further assumed to follow Poisson observation model  $x_t | N_t, \phi \sim \text{Poi}(\phi N_t)$ , where  $\phi$  is a scale parameter.

The goal is to compute the SL posterior for the parameters  $\boldsymbol{\theta} = (\log(r), \theta, K, \phi, \sigma_\varepsilon)$  given data  $\mathbf{x} = (x_t)_{t=1}^T$  with  $T = 100$ . The initial population size is  $N_0 = 1$  and we use independent priors  $\log(r) \sim \mathcal{U}([2, 5])$ ,  $\theta \sim \mathcal{U}([0.01, 2])$ ,  $K \sim \mathcal{U}([1, 5])$ ,  $\phi \sim \mathcal{U}([4, 20])$ ,  $\sigma_\varepsilon \sim \mathcal{U}([0, 0.8])$ . We use  $N = 100$  repeated simulations and the 13 summary statistics proposed by Wood (2010) to compute log-SL. We acknowledge that these statistics may be unideal as they are designed for the simpler Ricker model but this way we obtain a challenging target density which could be encountered during a typical LFI workflow. The “true” parameter for generating the data is  $\boldsymbol{\theta}_{\text{true}} = (3.5, 1.0, 3.5, 10, 0.3)$ . We use  $t_{\text{init}} = 20$  and  $i_{\text{MH}} = 2 \cdot 10^5$ . The initial location and initial proposal covariance are  $\boldsymbol{\theta}^{(0)} = (3.4, 0.9, 3.0, 8.0, 0.3)$  and

$\Sigma_0 = \text{diag}(0.05, 0.1, 0.25, 0.5, 0.05)^2$ , respectively. We observe noise level of  $\sigma_n \gtrsim 1.0$  in the highest posterior density region.

Figure 5 shows that EPoE produces the best sample-efficiency (especially when  $\varepsilon = 0.35$  or 0.3) though the improvements are not as substantial as in the more ideal GP modelling scenario of Section 5.1. We see that  $\varepsilon = 0.4$  does not produce accurate posterior (TV  $\gtrsim 0.3$ ) but already  $\varepsilon = 0.35$  produces reasonable approximations (TV  $\lesssim 0.1$ ) and requires only 300 – 500 log-SL evaluations. Naive with  $\varepsilon = 0.25$  would have required more evaluations than our limit  $10^3$  which we set to keep the computation time bounded. GP-MH and MH-BLFI again produce similar results.

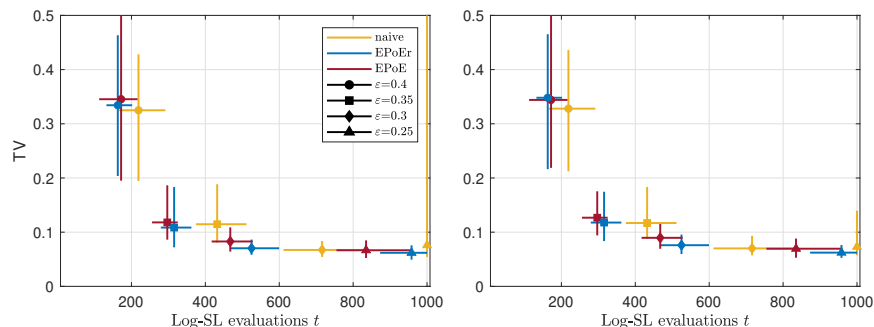


Figure 5: Accuracy of the marginal posterior approximation in the theta-Ricker experiment as a function of log-likelihood evaluations at the final iteration  $i_{\text{MH}} = 2 \cdot 10^5$ . *Left plot* shows GP-MH and the *right plot* the corresponding results by MH-BLFI.

Figure 6 shows a typical estimated posterior with EPoE and  $\varepsilon = 0.3$ . Parameters  $K$  and  $\phi$  cannot be fully identified which makes their true joint marginal posterior difficult to approximate also. The overall approximation quality is still good when  $\varepsilon \lesssim 0.35$  and, most importantly, the non-identifiability of these two parameters is clearly captured. When using  $\varepsilon = 0.4$  the approximation for  $(\log(r), \theta, \sigma_e)$  was still reasonable but substantial amount of the probability mass of  $(K, \phi)$  was often missed.

We were unable to obtain reasonable posterior approximations in our theta-Ricker experiment using BLFI. The main reason for this was that log-SL behaves irregularly on certain boundary regions of the parameter space where the method typically needs to evaluate. This produces a vicious cycle where a poor global GP fit causes next evaluation locations to be uninformative and the GP fit to remain poor. We unsuccessfully tried thresholding the log-SL values and special initialisations. Cropping the problematic parameter regions would help but is cumbersome due to the parameter correlations and would complicate optimising the acquisition functions. Similar problems emerge with other techniques relying on global GP surrogate. Especially current implementations of BOLFI typically excessively evaluate near the boundaries as seen e.g. in Järvenpää et al. (2019); Picchini et al. (2022). However, when initialised near the highest posterior density region, GP-MH and MH-BLFI produced accurate results and gracefully avoided the problematic regions. Namely, if a new parameter is proposed in such a region, a new log-likelihood evaluation is often triggered there. This point is then simply rejected without updating the GP or having to alter the acquisition function (see Appendix E.2 for details). In rare cases where the algorithm still

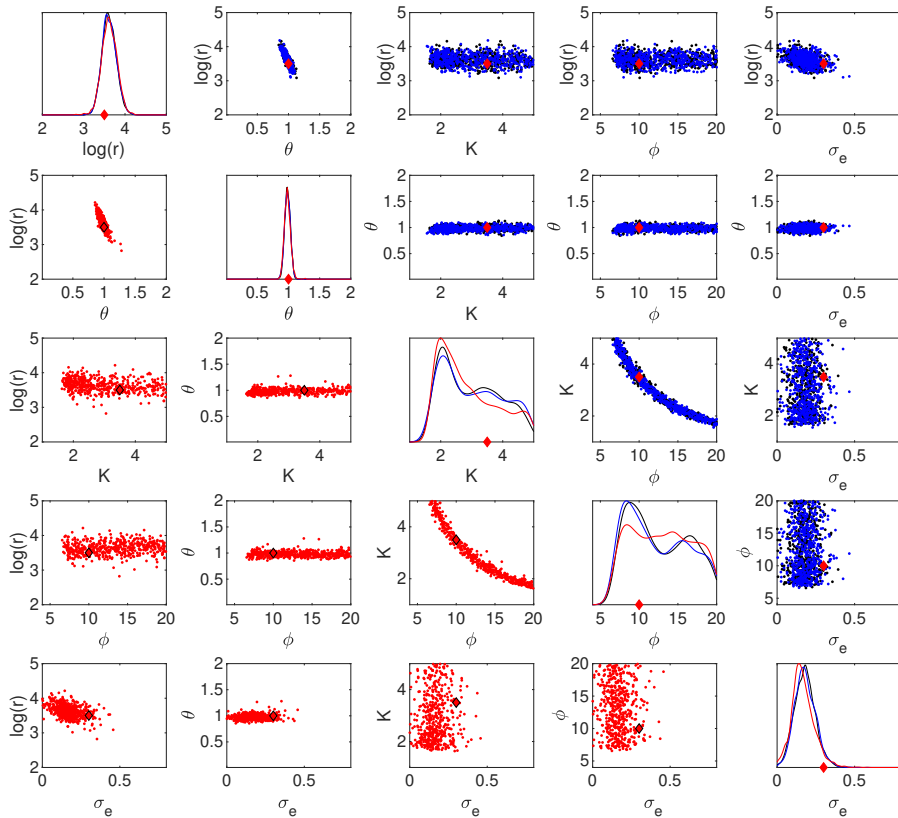


Figure 6: Comparison of the ground-truth posterior (red dots/line) computed using SL-MCMC and a typical example of estimated posterior (blue dots/line: GP-MH posterior, black dots/line: MH-BLFI posterior) in the case of theta-Ricker experiment. The red diamond shows the true parameter.

enters such problematic regions and gets stuck there, the algorithm can be restarted with different initialisation.

### 5.2.2 CELL BIOLOGY MODEL

We consider a model used for estimating cell motility and proliferation which are further needed when assessing the efficacy of certain medical treatments. Background and details of the model can be found in Price et al. (2018) and references therein. In the model T3T cells are represented as an  $R \times C$  binary matrix at each time point and each  $(x, y)$  location indicates whether a cell is present there or not. The cell dynamics are simulated over time using a random walk model which features two parameters that control the cell movement and reproduction, the probability of motility  $P_m \in [0, 1]$  and the probability of proliferation  $P_p \in [0, 1]$ . Observed data can be obtained using a scratch assay and consist of images (binary matrices), measured at some time points. The resulting likelihood function is intractable. We consider similar set-up as in Price et al. (2018). We use simulated data consisting of binary matrices over 145 time instances, we set  $R = 27, C = 36$  and we initially place 110 cells randomly in the rectangle with positions  $x \in \{1, 2, \dots, 13\}, y \in \{1, 2, \dots, 36\}$ .



Although the model has only two parameters, the imposed small budget for log-SL evaluations ( $\leq 750$ ) and the relatively large noise level of the log-likelihood evaluations (we use  $N = 2500$  which results  $\sigma_n \gtrsim 2.5$  in the highest density region) makes inference challenging. One log-SL evaluation takes approximately 7s (on a PC laptop with Intel Core i5 8265U and using the C-code by Price et al., 2018) which makes SL-MCMC feasible but fairly expensive. For example, a single chain with length  $10^5$  requires approximately 8 days of computing. Using more image data, larger lattice, more complicated cell dynamics or less efficient implementation would make the inference even more expensive. We use the same 145 summary statistics as Price et al. (2018). These include the Hamming distances between all the subsequent binary matrices over the 144 time intervals and the total number of cells in the final time period.

We use  $\theta_{\text{true}} = (0.35, 1.0 \cdot 10^{-3})$ . We first experimented with  $\mathcal{U}([0, 1]^2)$  prior but immediately observed that initialising our method in  $\theta_2 \gtrsim 4.0 \cdot 10^{-3}$ , where log-SL value is negligible and has very large variance, would not work. Similar difficulties would affect also SL-MCMC. In fact, log-SL decreases very fast near all boundaries which is problematic for B(O)LFI. While it is feasible to construct a bounding box to crop such regions in this particular 2D case, this in general involves tedious manual work. While not absolutely necessary, we restricted the parameter space of  $\theta_2$  and coded this into the prior  $\theta \sim \mathcal{U}([0, 1] \times [0, 4.0 \cdot 10^{-3}])$ . We use  $t_{\text{init}} = 10$ ,  $i_{\text{MH}} = 10^5$ ,  $\theta^{(0)} = (0.5, 1.5 \cdot 10^{-3})$  and  $\Sigma_0 = \text{diag}(0.02, 2.0 \cdot 10^{-4})^2$ .

Overall the results summarised as Figure 7 are similar to those in the previous experiments. However, the difference between the methods is smaller and the variability in the number of used log-SL evaluations is substantially larger especially when  $\varepsilon = 0.2$ . This variability is mostly caused by some individual iterations requiring around 30 – 70 evaluations. Presumably the fairly large  $\sigma_n$  and the low-dimensional  $\theta$  make the progress of the algorithm more dependent on randomness as in our other experiments. We nevertheless obtain reasonable posterior approximations with only a few hundred log-SL evaluations. The computational cost of each run was at most one to two hours. We observed isolated cases where the algorithm had traversed to the problematic boundary region and got stuck there. As these rare cases all happened when  $\varepsilon \geq 0.3$ , it is safe to say that our algorithm worked robustly in this experiment.




Figure 7: Accuracy of the posterior approximation in the Cell biology experiment as a function of log-likelihood evaluations at the final iteration  $i_{\text{MH}} = 10^5$ . *Left plot* shows GP-MH and the *right plot* the corresponding results by MH-BLFI.

### 5.3 Likelihood-free generalised Bayesian inference

We have now covered the case where one has access to a small number of noisy evaluations of the likelihood function or some approximation such as SL. However, GP-MH can be extended to other situations, in particular to the ABC case when a single model simulation is computationally costly and a surrogate model is formed for the ABC discrepancy. See Section A.2 for details on ABC relevant to what follows and Meeds and Welling (2014) for another approach that is also mentioned in Section 2.2.

Suppose the discrepancy  $\Delta(\mathbf{x}_o, \mathbf{x})$  follows Gaussian distribution with mean  $d(\boldsymbol{\theta})$  and variance  $\sigma_n^2(\boldsymbol{\theta})$  when  $\mathbf{x} \sim \pi(\mathbf{x} | \boldsymbol{\theta})$  and where the observed data  $\mathbf{x}_o$  is again suppressed from the notation. The ABC likelihood (A.4) of Appendix A.2 is then proportional to  $\Phi[(h - d(\boldsymbol{\theta}))/\sigma_n(\boldsymbol{\theta})]$  where  $h > 0$  is the threshold. The unknown mean discrepancy  $d(\boldsymbol{\theta})$  can be modelled with GP as first suggested by Gutmann and Corander (2016), though the likelihood ratio  $\Phi[(h - d(\boldsymbol{\theta}'))/\sigma_n(\boldsymbol{\theta}')]/\Phi[(h - d(\boldsymbol{\theta}))/\sigma_n(\boldsymbol{\theta})]$  in the MH accept/reject test then depends on  $d$  in a such way which does not seem to allow analytical computations. Although efficient numerical approach may be feasible, in the following we instead consider an alternative target density which directly facilitates GP-MH.

#### 5.3.1 GENERALISED POSTERIOR WITH EXPECTED ABC DISCREPANCY AS THE LOSS

We consider a target density proportional to

$$\pi(\boldsymbol{\theta})e^{-wl(\boldsymbol{\theta}, \mathbf{x}_o)}, \quad (30)$$

where  $w > 0$  is a scaling constant and the loss function  $l(\boldsymbol{\theta}, \mathbf{x}_o)$  is selected as

$$l(\boldsymbol{\theta}, \mathbf{x}_o) = l_{\Delta}(\boldsymbol{\theta}, \mathbf{x}_o) \triangleq \mathbb{E}_{\mathbf{x} | \boldsymbol{\theta}} \Delta(\mathbf{x}_o, \mathbf{x}). \quad (31)$$

We assume  $\mathbb{E}_{\mathbf{x} | \boldsymbol{\theta}} \Delta(\mathbf{x}_o, \mathbf{x})$  is finite at each  $\boldsymbol{\theta}$  and (30) can be normalised. This approach can be considered as an instance of generalised Bayesian inference by Bissiri et al. (2016) though the discrepancy  $\Delta(\mathbf{x}_o, \mathbf{x})$  in the ABC case typically does not satisfy a coherence property that Bissiri et al. (2016) used to provide a justification for the exponential form of (30). The density (30) is also a solution to

$$\arg \min_{\nu} \int wl(\boldsymbol{\theta}, \mathbf{x}_o)\nu(d\boldsymbol{\theta}) + \text{KL}(\nu || \pi). \quad (32)$$

For small  $w$  the generalised posterior (30) is similar to the prior  $\pi(\boldsymbol{\theta})$  in the sense of Kullback-Leibler divergence (KL). If  $w$  is set arbitrarily large, then a point mass at  $\arg \min_{\boldsymbol{\theta}} l(\boldsymbol{\theta}, \mathbf{x}_o)$  is a solution to (32) and can in principle be computed using (standard) Bayesian optimisation.

The choice  $l(\boldsymbol{\theta}, \mathbf{x}_o) = -\log \pi(\mathbf{x}_o | \boldsymbol{\theta})$  with  $w = 1$  gives the exact posterior and the ABC posterior is formally obtained when  $w = 1$  and

$$l(\mathbf{x}_o, \boldsymbol{\theta}) = l_{\text{ABC}}(\mathbf{x}_o, \boldsymbol{\theta}) \triangleq -\log \mathbb{E}_{\Delta | \boldsymbol{\theta}} K_h(\Delta) = -\log \int K_h(\Delta(\mathbf{x}_o, \mathbf{x}))\pi(\mathbf{x} | \boldsymbol{\theta}) d\mathbf{x}. \quad (33)$$

Assuming the kernel  $K_h$  is strictly positive and by using Jensen's inequality, we have  $l_{\text{ABC}}(\boldsymbol{\theta}, \mathbf{x}_o) \leq -\mathbb{E}_{\Delta | \boldsymbol{\theta}} \log K_h(\Delta)$ . When the Exponential kernel  $K_h(r) \propto e^{-r/h}$  is used, it is easy to see that  $-\mathbb{E}_{\Delta | \boldsymbol{\theta}} \log K_h(\Delta)$  equals  $wl_{\Delta}(\boldsymbol{\theta}, \mathbf{x}_o)$  with  $w = 1/h$  up to an additive

constant. That is, in this case (30) is also obtained from the ABC posterior by replacing the ‘‘ABC-loss’’  $l_{\text{ABC}}(\boldsymbol{\theta}, \mathbf{x}_o)$  with the larger loss  $wl_{\Delta}(\boldsymbol{\theta}, \mathbf{x}_o)$ .

A computational advantage of (30) is that the mean discrepancy only needs to be accurately estimated though meaningful uncertainty estimates are still needed to apply GP-MH. The standard ABC approach instead requires the more difficult task of estimating the lower tail probability of the discrepancy. Also, (30) may be more appropriate when the simulator-based model is supposedly grossly misspecified yet no better model is available. In this case the interpretation of the exact posterior is problematic. A downside is that (30) does not have the interpretation as the exact posterior when the model is well-specified and careful selection of  $w$  is required.

GP-MH is applied by replacing the target log-likelihood  $f(\boldsymbol{\theta})$  with  $-wl_{\Delta}(\boldsymbol{\theta}, \mathbf{x}_o)$ . The noisy evaluations are obtained as  $-w \sum_{i=1}^M \Delta(\mathbf{x}_o, \mathbf{x}^{(i)})/M$ , where  $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)}$  are samples from  $\pi(\mathbf{x} | \boldsymbol{\theta})$ , which clearly defines an unbiased estimator for  $-wl_{\Delta}(\boldsymbol{\theta}, \mathbf{x}_o)$ . Although the discrepancy is commonly formed to be non-negative, Gaussian distribution for  $\Delta(\mathbf{x}_o, \mathbf{x}^{(i)})$  is a sensible general assumption (Järvenpää et al., 2020). If a relatively large  $M$  can be used, the central limit theorem further justifies the Gaussianity assumption.

### 5.3.2 BACTERIAL COLONISATION DYNAMICS MODEL

Numminen et al. (2013) developed a simulator-based model to understand transmission dynamics of bacterial pathogen colonisation in day care centres, defined by combining a latent continuous-time Markov process with an cross-sectional observation model. The model has been used previously by Gutmann and Corander (2016); Järvenpää et al. (2020) to demonstrate that GP-based ABC methods can substantially improve the computational efficiency over standard ABC samplers though here we instead briefly investigate the loss-based general Bayesian inference with GP-MH as an alternative to ABC. We estimate the three unknown parameters of the model, the internal, external and co-infection parameters  $\beta, \Lambda$  and  $\theta$ , respectively. We use the same real data as Numminen et al. (2013) which describes colonisations with the bacterium *Streptococcus pneumoniae* and consists of varying numbers of sampled attendees at 29 day care centres at a single time point. Further details of the model and data can be found in Numminen et al. (2013).

The discrepancy is defined similarly as in Gutmann and Corander (2016). The data from each day care centre corresponds to a binary matrix whose elements specify whether each sampled attendee carries any of the circulating bacterial strains at the observed time point and is comprised using four summary statistics. The  $L^1$  distances between the empirical CDFs over the 29 day care centres both for the observed and simulated data are then computed for each summary. The discrepancy is defined as the square root of the average of the resulting four  $L^1$  distances (weighted by the magnitude of the observations). Also as in previous work, we use uniform priors  $\beta \sim \mathcal{U}([0, 11])$ ,  $\Lambda \sim \mathcal{U}([0, 2])$  and  $\theta \sim \mathcal{U}([0, 1])$ . We treat  $\sigma_n^2 = \mathbb{V}_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)} | \boldsymbol{\theta}}(-w \sum_{i=1}^M \Delta(\mathbf{x}_o, \mathbf{x}^{(i)})/M)$  as a constant with respect to  $\boldsymbol{\theta}$  and estimate it as in Section 5.1. We use  $t_{\text{init}} = 10$ ,  $i_{\text{MH}} = 10^5$ ,  $\boldsymbol{\theta}^{(0)} = (2.0, 0.4, 0.2)$  and  $\boldsymbol{\Sigma}_0 = \text{diag}(0.5, 0.1, 0.02)^2$  for each experiment.

Figure 8 shows the results with some choices of  $w$  and  $M$ . A few hundred ( $w = 10$  or  $20$ ) or a thousand ( $w = 40$ ) model simulations are sufficient to produce useful approximations. GP-MH is hence a feasible option also in this setting. Naive method interestingly works well

and clearly produces the best final accuracy when  $w = 40$  and  $M = 4$ . EPoEr and especially EPoE rely more on the GP model than naive and their performance is presumably more affected by the misspecified constant noise assumption of the GP. While naive requires more evaluations than the other methods on average to make the unconditional error smaller than  $\varepsilon$ , it appears that the extra evaluations are here not redundant and cause later accept/reject decisions to be made more accurately than required. The accuracy of the accept/reject decisions by EPoE and EPoEr are more closely centred near the upper bound  $\varepsilon$ .

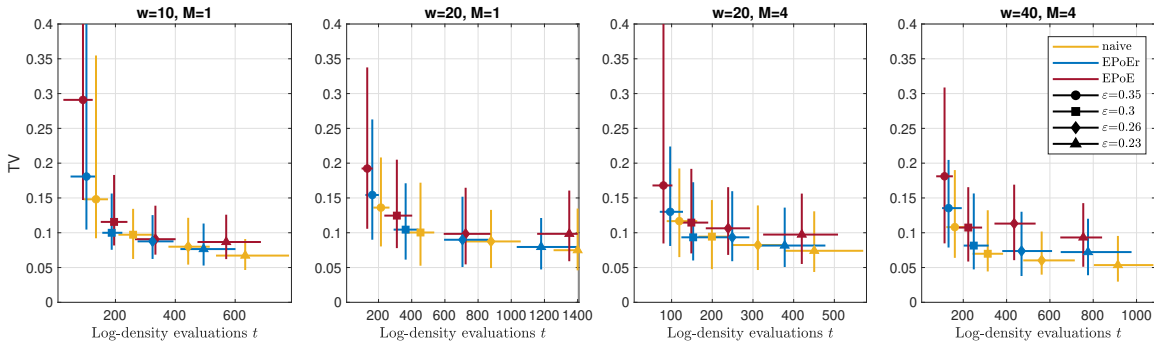


Figure 8: Accuracy of the marginal generalised posterior approximation of GP-MH in the bacterial infections experiment as a function of log-density evaluations at the final iteration  $i_{MH} = 10^5$ . The results for MH-BLFI were similar.

Some estimated generalised posteriors are illustrated in Figure 9 in the EPoEr case with  $\varepsilon = 0.23$ ,  $M = 1$  and  $w = 20$ . We observe that the choice  $w = 20$  (which is approximately equal to  $2(\mathbb{V}_{\mathbf{x}|\boldsymbol{\theta}}\Delta(\mathbf{x}_o, \mathbf{x}))^{-1/2}$  for any  $\boldsymbol{\theta}$  near the mode) produces fairly similar density estimate as the ABC posterior obtained using the ABC-MCMC sampler by Marjoram et al. (2003), the same discrepancy as for GP-MH and a uniform kernel with threshold  $h = 1.0$ . This suggests that, at least in some situations, our generalised posterior may be considered as an approximation to the ABC posterior itself.




Figure 9: Comparison of the ground-truth generalised posterior (red line) and 10 examples of the estimated posterior produced by GP-MH (blue lines) in the case of bacterial infections model. The ABC posterior (yellow line) is also shown for comparison.

## 6. Summary and discussion

We proposed a new sample-efficient framework for approximate Bayesian inference by combining Metropolis-Hastings sampling with Gaussian process emulation. The resulting GP-MH implementation is mainly intended for low- to moderate-dimensional problems ( $p < 10$ ) when only a small number ( $\sim 10^3$ ) of typically noisy likelihood evaluations (e.g.  $t < 10^3$ ) can be accessed. As is typically assumed in likelihood-free inference for simulator-based statistical models the underlying log-likelihood needs to be a smooth function and we further assume that the noise in evaluations is approximately Gaussian distributed within the region where the posterior distribution is concentrated. Among other analyses, we provided a Bayesian decision-theoretic justification for various parts of the method and discussed the relationship between GP-MH and “Bayesian optimisation-like” B(O)LFI frameworks for approximate Bayesian inference by Gutmann and Corander (2016); Järvenpää et al. (2021).

Two particular implementations of the proposed framework were investigated numerically, codes for reproducing the experiments and visualisations are available at <https://github.com/mjarvenpaa/GP-MH>. Our GP-MH and MH-BLFI implementations produced similar posterior approximations which is unsurprising given their connection (Section 4.2) and because most of the log-likelihood evaluations are typically collected during the early iterations which are neglected anyway as “burn-in”. The slightly changing target distribution of GP-MH thus does not appear problematic. Our experiments show that Bayesian sequential design strategies can provide worthwhile improvements to sample-efficiency over a naive strategy. However, we believe that their potential did not yet fully realise due to practical challenges with surrogate modelling. It is worth noting that the EPoE strategy does not account for the resulting posterior approximation directly (unlike IMIQR by Järvenpää et al., 2021), but is designed to minimise the number of evaluations needed to make the current MH accept/reject decision within the required accuracy. This suggests the more heuristic EPoEr as a pragmatic choice for GP-MH in applications.

GP-MH is based on the idea of building the GP surrogate model where the posterior distribution is concentrated, in order to make it more robust to aberrant behaviour of simulations near the parameter boundaries. This is in contrary to earlier B(O)LFI methods where the entire log-likelihood (or discrepancy function) is modelled with a single GP and the acquisition function optimised over the whole parameter space. On the other hand, posterior densities that do behave smoothly but feature many distant modes are better suited for alternative methods such as B(O)LFI. Computations needed to apply EPoE(r) methods are substantially more efficient than those developed by Järvenpää et al. (2021). For example, the GP computations in a typical run with the theta-Ricker model using  $2 \cdot 10^5$  iterations and around 500 evaluations took less than 15 minutes while the global optimisation of the IMIQR acquisition function only once in BLFI already takes up to one minute on the same computing settings. EPoEr and naive methods do not even need auxiliary optimisation which further simplifies their implementation in practice.

A potential downside over B(O)LFI is that GP-MH may require more careful initialisation. As with all MH samplers, selection of good proposal covariance can also be tricky. Also, if GP-MH is started far from the highest density region or if the log-likelihood behaves irregularly around the initial point, difficulties with either traversing to the highest density region or GP fitting might still emerge. The trade-off between accuracy and computational

cost needs thus to be adjusted, admittedly in a somewhat nontransparent fashion, using the parameter  $\varepsilon$ . A pragmatic advice would be to first run the algorithm using a fairly large  $\varepsilon$  while guided by the presented experimental results. Should the obtained posterior approximation appear inaccurate, or if less evaluations were used than anticipated, the algorithm could then be rerun with a decreased value of  $\varepsilon$  using the evaluations from the previous run as initial data.

There is still considerable room to further enhance and extend the proposed approach. Detailed analysis of the interplay between the error tolerance  $\varepsilon$ , the number of total evaluations needed and the accuracy of the resulting posterior approximation would be beneficial, albeit challenging future research topics. For global optimisation, it would be possible to consider extending GP-MH to work with simulated annealing. Use of GP-MH-like approach for the very challenging case of high-dimensional parameter space, improved modelling of the log-likelihood variance and alternative criteria for controlling the accuracy of the MH accept/reject test are additional topics that appear promising to pursue in the future.

## Acknowledgments

Some of the computations were performed on resources provided by UNINETT Sigma2 - the National Infrastructure for High Performance Computing and Data Storage in Norway. This research was funded by the Norwegian Research Council FRIPRO grant no. 299941, by the European Research Council grant no. 742158, and by Research Council of Norway through its Centre of Excellence Integreat - The Norwegian Centre for knowledge-driven machine learning, project number 332645.

## Appendix A. Background on likelihood-free inference

### A.1 Synthetic likelihood

LFI methods mainly differ on how the model simulations are used to approximate the intractable likelihood function. Synthetic likelihood (Wood, 2010; Price et al., 2018) is a parametric approximation to the likelihood  $\pi(\mathbf{x}_o | \boldsymbol{\theta})$  formed by first replacing the full data  $\mathbf{x}_o \in \mathbb{R}^d$  with summary statistics  $S(\mathbf{x}_o)$ , where  $S : \mathbb{R}^d \rightarrow \mathbb{R}^s$ ,  $s < d$ , and then assuming

$$\pi(S(\mathbf{x}_o) | \boldsymbol{\theta}) = \mathcal{N}_s(S(\mathbf{x}_o) | \boldsymbol{\mu}_\boldsymbol{\theta}, \boldsymbol{\Sigma}_\boldsymbol{\theta}). \quad (\text{A.1})$$

The approximation results from replacing the full data  $\mathbf{x}_o$  with potentially insufficient summary statistics  $S(\mathbf{x}_o)$  and from the Gaussianity assumption in (A.1) that rarely holds exactly. The unknown expectation  $\boldsymbol{\mu}_\boldsymbol{\theta} \in \mathbb{R}^s$  and covariance matrix  $\boldsymbol{\Sigma}_\boldsymbol{\theta} \in \mathbb{R}^{s \times s}$  in (A.1) are estimated from  $N$  repeated simulations for each proposed  $\boldsymbol{\theta}$  using

$$\hat{\boldsymbol{\mu}}_\boldsymbol{\theta} = \frac{1}{N} \sum_{i=1}^N S(\mathbf{x}_\boldsymbol{\theta}^{(i)}), \quad \hat{\boldsymbol{\Sigma}}_\boldsymbol{\theta} = \frac{1}{N-1} \sum_{i=1}^N (S(\mathbf{x}_\boldsymbol{\theta}^{(i)}) - \hat{\boldsymbol{\mu}}_\boldsymbol{\theta})(S(\mathbf{x}_\boldsymbol{\theta}^{(i)}) - \hat{\boldsymbol{\mu}}_\boldsymbol{\theta})^\top, \quad (\text{A.2})$$

where  $\mathbf{x}_\boldsymbol{\theta}^{(i)} \sim \pi(\cdot | \boldsymbol{\theta})$  for  $i = 1, \dots, N$ . An estimator for the likelihood function is then obtained by replacing the unknown  $\boldsymbol{\mu}_\boldsymbol{\theta}$  and  $\boldsymbol{\Sigma}_\boldsymbol{\theta}$  in (A.1) with the point estimates  $\hat{\boldsymbol{\mu}}_\boldsymbol{\theta}$  and  $\hat{\boldsymbol{\Sigma}}_\boldsymbol{\theta}$  in (A.2). See Section E.1 for some remarks on this estimator. The resulting log-synthetic

likelihood evaluations (abbreviated as log-SL) are noisy because  $N$  cannot in practice be large for computational reasons. Various extensions of SL have also been proposed (An et al., 2019, 2020; Frazier et al., 2022; Thomas et al., 2022) which similarly produce noisy log-likelihood approximations.

## A.2 Approximate Bayesian computation

The ABC likelihood is formed in a non-parametric fashion so that

$$\pi(\mathbf{x}_o | \boldsymbol{\theta}) \approx \pi_{\text{ABC}}(\mathbf{x}_o | \boldsymbol{\theta}) \triangleq \mathbb{E}_{\mathbf{x} | \boldsymbol{\theta}} K_h(\Delta(\mathbf{x}_o, \mathbf{x})) = \int K_h(\Delta(\mathbf{x}_o, \mathbf{x})) \pi(\mathbf{x} | \boldsymbol{\theta}) d\mathbf{x}, \quad (\text{A.3})$$

where the discrepancy  $\Delta(\mathbf{x}_o, \mathbf{x})$  quantifies the similarity between the observed data  $\mathbf{x}_o$  and the pseudo-data  $\mathbf{x}$  simulated from the model likelihood  $\pi(\mathbf{x} | \boldsymbol{\theta})$ . The data is commonly reduced to summary statistics (as in the SL case) so that, for example,  $\Delta(\mathbf{x}_o, \mathbf{x}) = \|S(\mathbf{x}_o) - S(\mathbf{x})\|$  where  $\|\cdot\|$  is some suitable norm although other choices are also possible. The kernel  $K_h(r)$  in (A.3) with bandwidth  $h > 0$  further controls the accuracy of the ABC approximation. A common choice is the uniform kernel  $K_h(r) \propto \mathbb{1}_{r \leq h}$  so that

$$\pi_{\text{ABC}}(\mathbf{x}_o | \boldsymbol{\theta}) \propto \mathbb{P}_{\mathbf{x} | \boldsymbol{\theta}}(\Delta(\mathbf{x}_o, \mathbf{x}) \leq h). \quad (\text{A.4})$$

Under some regularity conditions and assuming  $S$  above is a sufficient statistic, (A.4) converges to the exact likelihood as  $h \rightarrow 0$ . However, in practice, one usually needs to resort to non-sufficient statistics and  $h$  cannot be taken too small because simulation of pseudo-data that gives a close match to the observed summary statistic  $S(\mathbf{x})$  is inefficient.

An unbiased estimator for the ABC likelihood (A.3) is obtained as

$$\pi_{\text{ABC}}(\mathbf{x}_o | \boldsymbol{\theta}) \approx \frac{1}{M} \sum_{i=1}^M K_h(\Delta(\mathbf{x}_o, \mathbf{x}^{(i)})), \quad (\text{A.5})$$

where  $\mathbf{x}^{(i)} \sim \pi(\cdot | \boldsymbol{\theta})$  for  $i = 1, \dots, M$ . Further details on ABC methods can be found e.g. in Sisson et al. (2019).

## Appendix B. Proofs and mathematical derivations of Section 2.3 and 3

**Proof** [Proposition 1] We use the fact that the set of medians of a random variable is a closed interval which we denote as  $[m_1, m_2]$  where  $m_1 \leq m_2$ . That is,  $m$  is a median of  $\gamma \iff m \in [m_1, m_2]$ . From the definition of the median we see that  $\mathbb{P}(\gamma < m) \leq 1/2$  for any median  $m$ . Suppose  $m_1 < m \leq m_2$ . Since then  $1/2 \leq \mathbb{P}(\gamma \leq m_1) \leq \mathbb{P}(\gamma < m)$  it follows that  $\mathbb{P}(\gamma < m) = 1/2$  in this case.

We consider fixed  $u$ , treat the conditional error  $\mathcal{E}_{u, \hat{\gamma}}$  as a function of  $\hat{\gamma}$  and write it as

$$\mathcal{E}_{u, \hat{\gamma}} = \begin{cases} \mathbb{P}(\gamma < u) & \text{if } \hat{\gamma} \geq u, \\ \mathbb{P}(\gamma \geq u) & \text{if } \hat{\gamma} < u, \end{cases}$$

to ease up the analysis to follow. This function is clearly bounded and consists of two values depending whether  $\hat{\gamma} \geq u$  or  $\hat{\gamma} < u$ .

Suppose first  $u = m_1$ . Then  $\mathbb{P}(\gamma < m_1) \leq 1/2$  and consequently  $\mathbb{P}(\gamma \geq m_1) \geq 1/2$ . Thus, if we choose  $\hat{\gamma} \geq m_1$  we get the minimum. In particular, we can choose  $\hat{\gamma}$  to be any median.

Suppose now  $m_1 < u \leq m_2$ . Then  $\mathbb{P}(\gamma < u) = \mathbb{P}(\gamma \geq u) = 1/2$  so that any choice of  $\hat{\gamma}$  will do. Again, we can choose  $\hat{\gamma}$  to be any median.

Suppose  $u < m_1$ . Since  $u$  is not a median it must hold that  $\mathbb{P}(\gamma \geq u) < 1/2$  or  $\mathbb{P}(\gamma \leq u) < 1/2$ . The former option is not possible because it would contradict with the facts that CDF is a non-decreasing function and  $m_1$  is median. Thus  $\mathbb{P}(\gamma \leq u) < 1/2$  must hold and it clearly follows that  $\mathbb{P}(\gamma < u) < 1/2$  so we can choose  $\hat{\gamma}$  to be any median to get this minimum value. Similarly we see that if  $u > m_2$ , then  $\mathbb{P}(\gamma \geq u) < 1/2$  so that we can choose  $\hat{\gamma}$  to be any median to get the minimum. We have thus shown that the median of  $\gamma$  minimises the conditional error given each fixed value of  $u$ .

Since any median minimises the conditional error with each fixed  $u$ , it follows that any median minimises also the conditional error integrated over  $u \in [0, 1]$  which is the unconditional error. Alternatively, we can also see this as follows. We write

$$\begin{aligned} \mathcal{E}_{u, \hat{\gamma}} &= \mathbb{P}(\gamma < u | u) \mathbf{1}_{\hat{\gamma} \geq u} + \mathbb{P}(\gamma \geq u | u) \mathbf{1}_{\hat{\gamma} < u} \\ &= \mathbb{P}(\gamma < u | u) \mathbf{1}_{\hat{\gamma} \geq u} + (1 - \mathbb{P}(\gamma < u | u))(1 - \mathbf{1}_{\hat{\gamma} \geq u}) \\ &= \mathbf{1}_{\hat{\gamma} \geq u} (2\mathbb{P}(\gamma < u | u) - 1) + c_u, \end{aligned}$$

where  $c_u = \mathbb{P}(\gamma \geq u | u)$  does not depend on  $\hat{\gamma}$ . It follows that

$$E_{\hat{\gamma}} = \int_0^{\hat{\gamma}} (2\mathbb{P}(\gamma < u | u) - 1) du + \int_{\hat{\gamma}}^1 c_u du.$$

The claim then follows from the facts  $\mathbb{P}(\gamma < m_1) \leq 1/2$ ,  $\mathbb{P}(\gamma < m) = 1/2$  for  $m \in (m_1, m_2]$  and  $\mathbb{P}(\gamma < u) > 1/2$  for  $u > m_2$ .  $\blacksquare$

**Justification for (16) and (18).** The former equation is obtained as follows

$$\begin{aligned} \mathcal{E}_{t, u, \hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= \mathbb{P}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') < u | u) \mathbf{1}_{\text{med}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')) \geq u} + \mathbb{P}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq u | u) \mathbf{1}_{\text{med}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')) < u} \\ &= \mathbb{P}(\log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') < \tilde{u} | u) \mathbf{1}_{\text{med}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')) \geq u} + \mathbb{P}(\log \gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \tilde{u} | u) \mathbf{1}_{\text{med}(\gamma_f(\boldsymbol{\theta}, \boldsymbol{\theta}')) < u} \\ &= \Phi\left(\frac{\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u}}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}\right) \mathbf{1}_{\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') < \tilde{u}} + \Phi\left(\frac{\tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}\right) \mathbf{1}_{\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \tilde{u}} \\ &= \Phi\left(\frac{\min\{\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u}, \tilde{u} - \mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')\}}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}\right) \\ &= \Phi\left(-\frac{|\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u}|}{\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')}\right). \end{aligned}$$

We state a Lemma that we need to derive (18):

**Lemma B.1** *Suppose  $\sigma > 0$  and  $0 \leq a \leq b$ . Then*

$$\int_a^b \Phi\left(\frac{\log u - \mu}{\sigma}\right) du = e^\mu \left[ e^\beta \Phi\left(\frac{\beta}{\sigma}\right) - e^\alpha \Phi\left(\frac{\alpha}{\sigma}\right) + e^{\sigma^2/2} \left( \Phi\left(\frac{\alpha - \sigma^2}{\sigma}\right) - \Phi\left(\frac{\beta - \sigma^2}{\sigma}\right) \right) \right],$$

where  $\alpha \triangleq \log a - \mu$  and  $\beta \triangleq \log b - \mu$ .



**Proof** We first use change of variables  $x = (\log u - \mu)/\sigma$  to compute

$$\int_a^b \Phi\left(\frac{\log u - \mu}{\sigma}\right) du = \sigma e^\mu \int_{\alpha/\sigma}^{\beta/\sigma} e^{\sigma x} \Phi(x) dx.$$

The final equality then follows by using the equation 101.000 in Owen (1980) and some straightforward simplifications.  $\blacksquare$

To shorten the notation, we write  $\mu_t$  for  $\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  and  $\sigma_t$  for  $\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$ . We can write

$$E_{t,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \int_0^1 \Phi\left(-\frac{|\mu_t - \log u|}{\sigma_t}\right) du,$$

from which we see that if  $\mu_t \geq 0$ , then

$$E_{t,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \int_0^1 \Phi\left(\frac{\log u - \mu_t}{\sigma_t}\right) du.$$

The first case of (18) then follows immediately by using Lemma B.1 and some straightforward simplifications.

Suppose now that  $\mu_t < 0$ . Then  $0 < e^{\mu_t} < 1$  and we can write

$$\begin{aligned} E_{t,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= \int_0^{e^{\mu_t}} \Phi\left(\frac{\log u - \mu_t}{\sigma_t}\right) du + \int_{e^{\mu_t}}^1 \Phi\left(\frac{\mu_t - \log u}{\sigma_t}\right) du \\ &= 1 - e^{\mu_t} + \int_0^{e^{\mu_t}} \Phi\left(\frac{\log u - \mu_t}{\sigma_t}\right) du - \int_{e^{\mu_t}}^1 \Phi\left(\frac{\log u - \mu_t}{\sigma_t}\right) du. \end{aligned} \quad (\text{B.1})$$

We use Lemma B.1 to compute both integrals in (B.1) and after some straightforward computations we obtain the second case of (18).

**Proof** [Proposition 2] Based on the GP model, we have

$$\mathbf{y}^* | \boldsymbol{\theta}^*, \mathcal{D}_t \sim \mathcal{N}_b(m_t(\boldsymbol{\theta}^*), c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \boldsymbol{\Lambda}^*).$$

Then, by Lemma 5.1 in Järvenpää et al. (2021), it follows that<sup>2</sup>

$$\begin{aligned} \begin{bmatrix} m_{t+b}^*(\boldsymbol{\theta}) \\ m_{t+b}^*(\boldsymbol{\theta}') \end{bmatrix} | \boldsymbol{\theta}^*, \mathcal{D}_t &\sim \mathcal{N}_2\left(\begin{bmatrix} m_t(\boldsymbol{\theta}) \\ m_t(\boldsymbol{\theta}') \end{bmatrix}, \begin{bmatrix} \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) & \omega_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) \\ \omega_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) & \tau_t^2(\boldsymbol{\theta}'; \boldsymbol{\theta}^*) \end{bmatrix}\right), \\ c_{t+b}^*(\boldsymbol{\theta}, \boldsymbol{\theta}') &= c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \omega_t(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*), \end{aligned}$$

where  $*$  is used to emphasise that these quantities depend on  $\boldsymbol{\theta}^*$  and possibly also  $\mathbf{y}^*$  via  $\mathcal{D}^*$ . It follows that

$$\mu_{t+b}^*(\boldsymbol{\theta}, \boldsymbol{\theta}') | \boldsymbol{\theta}^*, \mathcal{D}_t \sim \mathcal{N}_1(\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}'), \xi_t^2(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*)),$$

---

2. Lemma 5.1 in fact shows the result for the GP mean and variance functions in a single  $\boldsymbol{\theta}$ -location only but it is easy to see that the result immediately extends for the more general case considered here.

where  $\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is given by (13) and  $\xi_t^2(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*)$  by (22). We also see that

$$\begin{aligned}\sigma_{t+b}^{2*}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= s_t^2(\boldsymbol{\theta}) - \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) + s_t^2(\boldsymbol{\theta}') - \tau_t^2(\boldsymbol{\theta}'; \boldsymbol{\theta}^*) - 2(c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - \omega_t^2(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*)) \\ &= s_t^2(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta}') - 2c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') - (\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) + \tau_t^2(\boldsymbol{\theta}'; \boldsymbol{\theta}^*) - \omega_t^2(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*)) \quad (\text{B.2}) \\ &= \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*).\end{aligned}$$

To shorten the notation, we once again drop “ $(\boldsymbol{\theta}, \boldsymbol{\theta}')$ ” from various formulas. For example, we write  $\mu_t$  for  $\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  and  $\xi_t(\boldsymbol{\theta}^*)$  for  $\xi_t(\boldsymbol{\theta}', \boldsymbol{\theta}; \boldsymbol{\theta}^*)$ .

We write the conditional error as

$$\mathcal{E}_{t+b, u, \hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = 2\Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right) \mathbb{1}_{\mu_{t+b}^* \geq \tilde{u}} - \mathbb{1}_{\mu_{t+b}^* \geq \tilde{u}} + 1 - \Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right). \quad (\text{B.3})$$

We then compute

$$\begin{aligned}\mathbb{E}_{\mu_{t+b}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} \left[ \mathbb{1}_{\mu_{t+b}^* \geq \tilde{u}} \right] &= 1 - \int_{-\infty}^{\tilde{u}} \mathcal{N}(\mu_{t+b}^* | \mu_t, \xi_t^2(\boldsymbol{\theta}^*)) = 1 - \Phi\left(\frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)}\right), \\ \mathbb{E}_{\mu_{t+b}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} \left[ 1 - \Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right) \right] &= 1 - \Phi\left(\frac{\tilde{u} - \mu_t}{\sigma_t}\right),\end{aligned}$$

where we have used equation 3.82 in Rasmussen and Williams (2006) and the fact  $\Phi(z) = 1 - \Phi(-z)$ . The first term in (B.3) requires some more work. We write

$$\begin{aligned}\mathbb{E}_{\mu_{t+b}^* | \boldsymbol{\theta}^*, \mathcal{D}_t} \left[ 2\Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right) \mathbb{1}_{\mu_{t+b}^* \geq \tilde{u}} \right] &= 2 \int_{\tilde{u}}^{\infty} \Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right) \mathcal{N}(\mu_{t+b}^* | \mu_t, \xi_t^2(\boldsymbol{\theta}^*)) d\mu_{t+b}^* \\ &= 2\Phi\left(\frac{\tilde{u} - \mu_t}{\sigma_t}\right) - 2 \int_{-\infty}^{\tilde{u}} \Phi\left(\frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*}\right) \mathcal{N}(\mu_{t+b}^* | \mu_t, \xi_t^2(\boldsymbol{\theta}^*)) d\mu_{t+b}^* \\ &= 2\Phi\left(\frac{\tilde{u} - \mu_t}{\sigma_t}\right) - 2 \int_{-\infty}^{\frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)}} \Phi\left(\frac{\tilde{u} - \mu_t - \xi_t(\boldsymbol{\theta}^*)x}{\sigma_{t+b}^*}\right) \mathcal{N}(x | 0, 1) dx, \quad (\text{B.4})\end{aligned}$$

where we used the transformation  $x = (\mu_{t+b}^* - \mu_t)/\xi_t(\boldsymbol{\theta}^*)$ . To compute the integral in (B.4) we use the equation 10.010.1 in Owen (1980). After some straightforward computations we obtain

$$\int_{-\infty}^{\frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)}} \Phi\left(\frac{\tilde{u} - \mu_t - \xi_t(\boldsymbol{\theta}^*)x}{\sigma_{t+b}^*}\right) \mathcal{N}(x | 0, 1) dx = \text{BvN}\left(\frac{\tilde{u} - \mu_t}{\sigma_t}, \frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)}; \frac{\xi_t(\boldsymbol{\theta}^*)}{\sigma_t}\right), \quad (\text{B.5})$$

where  $\text{BvN}(h, k, \rho)$  is the pdf of a bivariate Gaussian with unit variances and correlation coefficient  $\rho$  evaluated at  $(h, k)^\top$ . We use the connection between BvN and Owen’s T function given by equation 3.1 in Owen (1980) (the first case of which applies here because

$hk = (\tilde{u} - \mu_t)^2 / (\sigma_t \xi_t(\boldsymbol{\theta}^*)) \geq 0$  and because  $hk = 0 \iff h = k = 0$  hold with (B.5)) and the fact  $T(h, 0) = 0$  for any  $h \in \mathbb{R}$ , to further obtain

$$\begin{aligned} & \text{BvN} \left( \frac{\tilde{u} - \mu_t}{\sigma_t}, \frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)}, \frac{\xi_t(\boldsymbol{\theta}^*)}{\sigma_t} \right) \\ &= \frac{1}{2} \Phi \left( \frac{\tilde{u} - \mu_t}{\sigma_t} \right) + \frac{1}{2} \Phi \left( \frac{\tilde{u} - \mu_t}{\xi_t(\boldsymbol{\theta}^*)} \right) - T \left( \frac{\tilde{u} - \mu_t}{\sigma_t}, \frac{\sqrt{\sigma_t^2 - \xi_t^2(\boldsymbol{\theta}^*)}}{\xi_t(\boldsymbol{\theta}^*)} \right). \end{aligned}$$

Once we combine the equations above, we see that all the  $\Phi(\cdot)$ -terms cancel out and we are left with (21).

The formula (20) follows immediately from above because we can change the order of integration over  $u \in [0, 1]$  and the expectation with respect to  $\pi(\mu_{t+b}^* | \mu_t, \xi_t(\boldsymbol{\theta}^*))$  using Fubini's theorem.

By using (19) and the fact  $\Phi(z) = 1 - \Phi(-z)$ , we write the expected variance of  $\kappa_{u,f}$  as

$$L_t^{v,u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = \int_{-\infty}^{\infty} \left[ \Phi \left( \frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*} \right) - \Phi^2 \left( \frac{\tilde{u} - \mu_{t+b}^*}{\sigma_{t+b}^*} \right) \right] \mathcal{N}(\mu_{t+b}^* | \mu_t, \xi_t^2(\boldsymbol{\theta}^*)) d\mu_{t+b}^*.$$

We then recognise that this integral is of the same form as in the proof of Lemma 3.1 in Järvenpää et al. (2019) from which (21) follows.  $\blacksquare$

**Proof** [Proposition 3] The result for the variance of the log-MH ratio follows immediately from  $\sigma_{t+b}^2(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  which is given by (B.2). The Owen's T function satisfies  $T(h, a) = \frac{1}{2\pi} \int_0^a \frac{e^{-h^2(1+x^2)/2}}{1+x^2} dx$ . This fact shows that the function  $a \mapsto T(h, \sqrt{a}), a \geq 0$  is strictly increasing with any fixed  $h \in \mathbb{R}$ . It follows that  $L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  is minimised when  $\frac{\sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}') - \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)}{\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)}$  is minimised which clearly happens when  $\boldsymbol{\theta}^*$  is chosen as in (25) since  $0 \leq \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) \leq \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}')$ . (These two inequalities follows from (C.1) and the fact that variance is always non-negative.) This reasoning holds with any  $u > 0$  so that  $L_t^E(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  is also minimised by this choice of  $\boldsymbol{\theta}^*$ . The proof for the case of  $L_t^{v,u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  is similar as for  $L_t^{\mathcal{E},u}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$ .  $\blacksquare$

## Appendix C. Some theoretical analysis

In the following we present some theoretical analysis for understanding GP-MH. The key results are briefly summarised in the main text as Section 4.4.

### C.1 Evaluation locations

We analyse the optimal evaluations locations. We mostly limit our attention to the sequential case  $b = 1$  where  $\boldsymbol{\theta}^*$  consists of a single parameter. Throughout this section we assume  $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$  are arbitrary distinct points.

C.1.1 ANALYSIS OF  $\tau_t$ -FUNCTION

We start by stating some results regarding  $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$  in (23) which gives the reduction of GP variance at  $\boldsymbol{\theta}$  resulting from evaluations at  $\boldsymbol{\theta}^*$ . When  $b = 1$  we can write (23) as

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \frac{c_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}^*)}{s_t^2(\boldsymbol{\theta}^*) + \sigma_n^2(\boldsymbol{\theta}^*)}.$$

Consider first the case  $\sigma_n(\boldsymbol{\theta}) = 0$ . We see immediately that then  $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}) = s_t^2(\boldsymbol{\theta})$  so that the intuitive choice  $\boldsymbol{\theta}^* = \boldsymbol{\theta}$  gives the maximal reduction of variance at  $\boldsymbol{\theta}$ .

Consider now the situation  $\sigma_n(\boldsymbol{\theta}) = \sigma_n > 0$  for all  $\boldsymbol{\theta}$ , that is, the noise level is constant. Then we have  $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}) = [s_t^2(\boldsymbol{\theta}) / (s_t^2(\boldsymbol{\theta}) + \sigma_n^2)] s_t^2(\boldsymbol{\theta}) < s_t^2(\boldsymbol{\theta})$  whenever  $s_t(\boldsymbol{\theta}) > 0$ . We can write  $c_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) = \rho_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) s_t(\boldsymbol{\theta}) s_t(\boldsymbol{\theta}^*)$ , where  $\rho_t(\cdot, \cdot)$  is the GP correlation function and

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \frac{\rho_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}^*) s_t^2(\boldsymbol{\theta}) s_t^2(\boldsymbol{\theta}^*)}{s_t^2(\boldsymbol{\theta}^*) + \sigma_n^2}.$$

If  $|\rho_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*)| = 1$  and  $s_t^2(\boldsymbol{\theta}^*) > s_t^2(\boldsymbol{\theta}) > 0$  then we obtain

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \frac{s_t^2(\boldsymbol{\theta}) s_t^2(\boldsymbol{\theta}^*)}{s_t^2(\boldsymbol{\theta}^*) + \sigma_n^2} = \frac{s_t^4(\boldsymbol{\theta})}{s_t^2(\boldsymbol{\theta}) + \frac{s_t^2(\boldsymbol{\theta})}{s_t^2(\boldsymbol{\theta}^*)} \sigma_n^2} > \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta})$$

which shows that choosing  $\boldsymbol{\theta}^* = \boldsymbol{\theta}$  does not produce the maximal reduction of variance at  $\boldsymbol{\theta}$  in general. This is obviously also the case when  $\sigma_n(\boldsymbol{\theta}) > 0$  is not constant with respect to  $\boldsymbol{\theta}$ . For example, if  $0 < s_t(\boldsymbol{\theta}) < \infty$  and  $\sigma_n(\boldsymbol{\theta}) = \infty$  then obviously  $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}) = 0$  but it is possible that  $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) > 0$  for some  $\boldsymbol{\theta}^* \neq \boldsymbol{\theta}$ .

C.1.2 ANALYSIS OF  $\xi_t$ -FUNCTION

The preliminary analysis above showed that an evaluation at  $\boldsymbol{\theta}$  may not maximally reduce the uncertainty about  $f$  at  $\boldsymbol{\theta}$  (or be a sensible choice at all) unless the evaluation is exact. We now similarly analyse  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  in (22). To this end, we first observe that we can alternatively write (22) as

$$\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = (c_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) - c_t(\boldsymbol{\theta}', \boldsymbol{\theta}^*)) [c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \boldsymbol{\Lambda}^*]^{-1} (c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}) - c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}')) \geq 0. \quad (\text{C.1})$$

First, suppose  $\sigma_n(\boldsymbol{\theta}) = \sigma_n(\boldsymbol{\theta}') = 0$ . Then  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; [\boldsymbol{\theta}, \boldsymbol{\theta}']) = s_t^2(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta}') - 2c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}')$  which is easily verified using (C.1) and some straightforward computations. That is, when two evaluations are to be used in the non-noisy setting so that  $b = 2$ , the intuitive choice  $\boldsymbol{\theta}^* = [\boldsymbol{\theta}, \boldsymbol{\theta}']$  produces the maximal reduction of uncertainty. Similarly, if we have already evaluated at  $\boldsymbol{\theta}$  and have one evaluation to use so that  $b = 1$ , then  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = \tau_t^2(\boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  which shows that the intuitive choice  $\boldsymbol{\theta}^* = \boldsymbol{\theta}'$  produces the maximal reduction of uncertainty which in this case equals  $s_t^2(\boldsymbol{\theta}') = \sigma_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}')$ .

Let us now suppose  $b = 1$ . We can then write (C.1) as

$$\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = \frac{[c_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) - c_t(\boldsymbol{\theta}', \boldsymbol{\theta}^*)]^2}{s_t^2(\boldsymbol{\theta}^*) + \sigma_n^2(\boldsymbol{\theta}^*)}. \quad (\text{C.2})$$

We suppose  $\sigma_n(\boldsymbol{\theta}) = \sigma_n \geq 0$  for all  $\boldsymbol{\theta}$ . As analysing (C.2) analytically is hard in general when  $b = 1$ , we limit our attention to the  $t = 0$  case where the GP posterior equals the GP prior. For simplicity, we also assume a stationary covariance function of the form  $k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_s^2 \kappa(\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_{\Lambda})$  where  $\sigma_s > 0$ ,  $\kappa : \mathbb{R}_+ \rightarrow [0, 1]$  is a strictly decreasing function with  $\kappa(0) = 1$  and  $\lim_{r \rightarrow \infty} \kappa(r) = 0$ ,  $\Lambda$  is a positive definite matrix and  $\|\boldsymbol{\theta}\|_{\Lambda}^2 \triangleq \boldsymbol{\theta}^{\top} \Lambda \boldsymbol{\theta}$ . For example, the choice  $\Lambda = \text{diag}(l_1^2, \dots, l_p^2)^{-1}$  where  $l_i > 0$  are the length-scales and  $\kappa(r) = \exp(-r^2/2)$  gives the (anisotropic) squared exponential (SE) covariance function. Then (C.2) can be written as

$$\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) = c[\kappa(\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_{\Lambda}) - \kappa(\|\boldsymbol{\theta}' - \boldsymbol{\theta}^*\|_{\Lambda})]^2, \quad (\text{C.3})$$

where  $c > 0$  does not depend on  $\boldsymbol{\theta}^*$ . Now, (C.3) shows that if  $\boldsymbol{\theta}^*$  is far enough (as measured using the norm  $\|\cdot\|_{\Lambda}$ ) from both  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$ , then  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*) \approx 0$  so that the uncertainty associated with the MH accept/reject decision will decrease only little. Surprisingly, this is also the case if  $\kappa(\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_{\Lambda}) = \kappa(\|\boldsymbol{\theta}' - \boldsymbol{\theta}^*\|_{\Lambda})$ , which is equivalent to  $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_{\Lambda} = \|\boldsymbol{\theta}' - \boldsymbol{\theta}^*\|_{\Lambda}$ . That is, points  $\boldsymbol{\theta}^*$  that are equally far from  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$  are uninformative in this case.

Under the above assumptions, we can reason that if  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$  are very far from each other, then choosing  $\boldsymbol{\theta}^* = \boldsymbol{\theta}$  or  $\boldsymbol{\theta}^* = \boldsymbol{\theta}'$  will approximately maximise  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$ . Suppose next that we have SE covariance function. Then  $\nabla_{\boldsymbol{\theta}^*} \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  is proportional to

$$\left[ e^{-\frac{1}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_{\Lambda}^2} - e^{-\frac{1}{2}\|\boldsymbol{\theta}' - \boldsymbol{\theta}^*\|_{\Lambda}^2} \right] \left[ e^{-\frac{1}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_{\Lambda}^2} \Lambda(\boldsymbol{\theta} - \boldsymbol{\theta}^*) - e^{-\frac{1}{2}\|\boldsymbol{\theta}' - \boldsymbol{\theta}^*\|_{\Lambda}^2} \Lambda(\boldsymbol{\theta}' - \boldsymbol{\theta}^*) \right].$$

This formula shows that the gradient  $\nabla_{\boldsymbol{\theta}^*} \xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  is nonzero both at  $\boldsymbol{\theta}^* = \boldsymbol{\theta}$  and  $\boldsymbol{\theta}^* = \boldsymbol{\theta}'$  because  $\boldsymbol{\theta} \neq \boldsymbol{\theta}'$  so these points are not local or global optima (except possibly in the special case where they belong to the boundary of the set  $\Theta$ ). Our numerical experiments in Section 5 further demonstrate that it matters in practice whether we optimise  $\xi_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\theta}^*)$  globally or over  $\boldsymbol{\theta}^* \in \{\boldsymbol{\theta}, \boldsymbol{\theta}'\}$ .

## C.2 Worst case upper bounds for the repeated, noisy evaluations

We provide some analysis on the number of repeated, noisy log-likelihood evaluations, denoted by  $b$  in this section, needed at an individual iteration of Algorithm 3 to satisfy the condition  $\mathcal{E}_{b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \leq \varepsilon$  or  $E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \leq \varepsilon$ . We first recognise some special cases: If  $\varepsilon \geq 1/2$ , then obviously no log-likelihood evaluations are needed. On the other hand, if  $\varepsilon = 0$  and  $\pi(\boldsymbol{\theta}') > 0$  then the log-likelihood must be known exactly at both  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$  which requires arbitrarily many evaluations in the noisy case. As seen in Figure 1, if  $\mu_t(\boldsymbol{\theta}, \boldsymbol{\theta}') = \log(u)$  then  $\mathcal{E}_{t,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \Phi(0) = 1/2$  even if  $\sigma_t(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is negligible yet nonzero. This situation however happens with vanishing probability.

**Proposition C.1** *Suppose that the GP prior model of Section 2.5 but with a stationary covariance function as in Section C.1.2 holds and that  $b \geq 0$  evaluation locations are then chosen as to minimise either the conditional error (C.4) or unconditional error (C.5) computed between distinct parameters  $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$  in an optimal fashion of Section 3. Suppose  $0 < \varepsilon < 1/2$  and  $\sigma_n(\boldsymbol{\theta}) > 0$  for all  $\boldsymbol{\theta} \in \Theta$ . Then<sup>3</sup> it holds that*

$$\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \varepsilon) \leq 1 - e^{2\Phi^{-1}(\varepsilon)c_b}, \quad (\text{C.4})$$

3. The fact  $c_b = \min\{2\sigma_s, 2\sqrt{2}\bar{\sigma}_n/\sqrt{b}\} \leq 2\sqrt{2}\bar{\sigma}_n/\sqrt{b}$  could be used to slightly simplify the bounds when  $b$  is even and  $b \geq 2$ .

$$E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \leq \max_{\mu \leq 0} \left\{ \Phi\left(\frac{\mu}{c_b}\right) + e^{\mu+c_b^2/2} \left( \Phi\left(-\frac{\mu+c_b^2}{c_b}\right) - 2\Phi(-c_b) \right) \right\}, \quad (\text{C.5})$$

where  $\mathbb{P}(\cdot)$  is with respect to  $u \sim \mathcal{U}([0,1])$ ,  $c_b \triangleq 2 \min\{\sigma_s, \bar{\sigma}_n/\sqrt{[b/2]}\}$  and  $\bar{\sigma}_n \triangleq \max\{\sigma_n(\boldsymbol{\theta}), \sigma_n(\boldsymbol{\theta}')\}$ .

Importantly, Proposition C.1 concerns the worst case situation where no log-likelihood evaluations are yet obtained (that is, the GP posterior equals the GP prior) and where the evaluations are always such that  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') = \Phi^{-1}(\varepsilon)\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')$  as seen in the proof of (C.4) in Appendix C.3. Furthermore, Proposition C.1 is for the batch case where  $b$  evaluations are selected simultaneously (but we hypothesise similar result holds in the sequential case). We observe numerically that the maximum of (C.5) is typically found in  $\mu \in (-0.7, 0)$ . We see that  $\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) \rightarrow 0$  and  $E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \rightarrow 0$  as  $b \rightarrow \infty$ .

We next develop revised bounds that are more representative of practice though they still assume a certain worst case situation. We consider a Gaussian target density  $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  where  $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times p}$  is positive definite and w.l.o.g. we set  $\boldsymbol{\mu} = \mathbf{0}$ . We consider a Gaussian proposal  $q(\boldsymbol{\theta}' | \boldsymbol{\theta}) = \mathcal{N}_p(\boldsymbol{\theta}' | \boldsymbol{\theta}, s^2\boldsymbol{\Sigma})$  where  $s > 0$  is a fixed scaling parameter. We suppose that the artificial scenario, where  $\boldsymbol{\theta}$  is first drawn from the target and  $\boldsymbol{\theta}'$  is then drawn from the proposal so that  $\boldsymbol{\theta} \sim \mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$ ,  $\boldsymbol{\theta}' | \boldsymbol{\theta} \sim \mathcal{N}_p(\boldsymbol{\theta}, s^2\boldsymbol{\Sigma})$ , represents a typical MH iteration. Then we obtain

$$\mathbb{E}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) = -\frac{1}{2}ps^2, \quad \mathbb{V}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) = \frac{1}{2}ps^2(s^2 + 2), \quad (\text{C.6})$$

where  $f$  is the Gaussian log-target density and where the expectation and variance are wrt. the randomness due to sampling  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$ . In particular, if we use the common recommended choice  $s^2 = 2.4^2/p$  (Gelman et al., 2013, Chapter 12.2), we obtain  $\mathbb{E}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) = -2.88$  and a relatively large standard deviation  $\text{sd}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) = 2.4\sqrt{2.88/p + 1}$ .

The bounds of Proposition C.1 are then revised by using the distribution of  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}')$ , which we assume to be the same as that of  $f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})$  described above, in place of its worst case choice. These computations are not completely analytic but the resulting bounds are easily evaluated numerically. The details are postponed to Appendix C.3.

Figure C.1 demonstrates the bounds. According to the bounds of Proposition C.1 (dashed lines in (a) and (b)) hundreds of repeated evaluations are needed if  $\varepsilon$  is small. As expected, the revised bounds (solid lines in (a) and (c)), based on the assumed distribution of  $\mu_t$  in the Gaussian case, are tighter. These results indicate a potential shortcoming of GP-MH: Many repeated evaluations may be needed at least in some individual worst case iterations of the GP-MH algorithm.

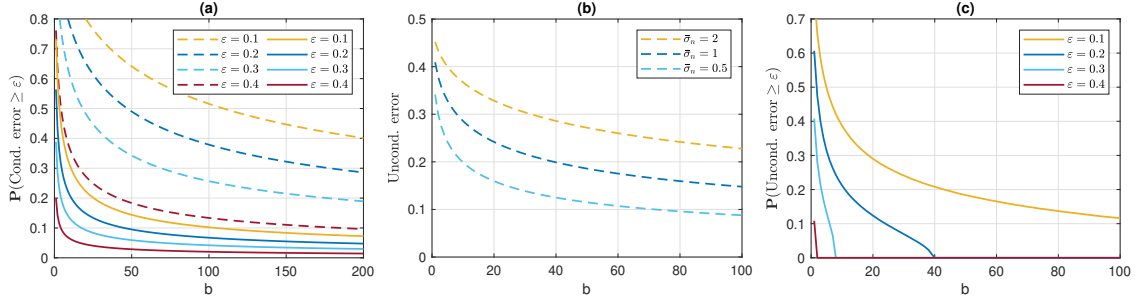


Figure C.1: (a) Dashed lines show the worst case bound (C.4) for the probability that the conditional error is larger than  $\varepsilon$ . Solid lines show the corresponding bound based on the assumed distribution of  $\mu_b$  in the Gaussian case (C.15) with  $p = 5$  and  $s^2 = 2.4^2/p$ . In both cases  $\bar{\sigma}_n = 1$ . (b) The worst case bound (C.5) for the unconditional error for three different values of  $\bar{\sigma}_n$ . (c) The upper bounds for the probability that the unconditional error is larger than  $\varepsilon$  in the assumed Gaussian case (C.16).

### C.3 Proofs and mathematical derivations of Section C.2

**Proof** [Proposition C.1] We denote  $\lambda_b \triangleq -\Phi^{-1}(\varepsilon)\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')$ . Note that  $\lambda_b > 0$  since  $\varepsilon \in (0, 1/2)$ . We then obtain

$$\begin{aligned}
 \mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) &= \mathbb{P}\left(\Phi\left(-\frac{|\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u}|}{\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')}\right) \geq \varepsilon\right) \\
 &= \mathbb{P}(|\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u}| \leq -\Phi^{-1}(\varepsilon)\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')) \\
 &= 1 - \mathbb{P}(\{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u} \geq \lambda_b\} \cup \{\tilde{u} - \mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \lambda_b\}) \\
 &= 1 - \mathbb{P}(\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \tilde{u} \geq \lambda_b) - \mathbb{P}(\tilde{u} - \mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \lambda_b) \\
 &= 1 - \mathbb{P}(u \leq e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b}) - \mathbb{P}(u > e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') + \lambda_b}) \\
 &= 1 - \min\{e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b}, 1\} - (1 - \min\{e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') + \lambda_b}, 1\}) \\
 &= \max\{1 - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b}, 0\} + \min\{e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') + \lambda_b} - 1, 0\}.
 \end{aligned} \tag{C.7}$$

If  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq 0$ , then simple computation shows that

$$\begin{aligned}
 \mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) &= \max\{1 - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b}, 0\} \\
 &= \begin{cases} 0 & \text{if } \lambda_b \leq \mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}'), \\ 1 - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b} & \text{otherwise.} \end{cases}
 \end{aligned}$$

Since the function  $\mu_b \mapsto 1 - e^{\mu_b - \lambda_b}$  is decreasing for  $\mu_b \geq 0$  and because  $\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) = 0$  for  $\lambda_b \leq \mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}')$ , we see that  $\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) \leq 1 - e^{-\lambda_b}$  for any  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq 0$ .

If  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') < 0$ , then we see that

$$\begin{aligned}
 \mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) &= 1 - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b} + \min\{e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') + \lambda_b} - 1, 0\} \\
 &= \begin{cases} 1 - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b} & \text{if } -\lambda_b \leq \mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}'), \\ e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') + \lambda_b} - e^{\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') - \lambda_b} & \text{otherwise.} \end{cases}
 \end{aligned}$$

Since the function  $\mu_b \mapsto 1 - e^{\mu_b - \lambda_b}$  is decreasing, its maximum in  $-\lambda_b \leq \mu_b < 0$  occurs with  $\mu_b = -\lambda_b$ . As  $\mu_b \mapsto e^{\mu_b - \lambda_b} e^{\mu_b + \lambda_b} - e^{\mu_b - \lambda_b} = 2e^{\mu_b} \sinh(\lambda_b)$  is increasing in  $\mu_b < -\lambda_b$ , the choice  $\mu_b = -\lambda_b$  gives an upper bound also when  $\mu_b < -\lambda_b$ . We have thus shown

$$\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}} \geq \varepsilon) \leq 1 - e^{-2\lambda_b} = 1 - e^{2\Phi^{-1}(\varepsilon)\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')}, \quad (\text{C.8})$$

which also works when  $\mu_b(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq 0$ .

Next we obtain

$$\begin{aligned} \sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}') &= \sqrt{s_b^2(\boldsymbol{\theta}) + s_b^2(\boldsymbol{\theta}') - 2c_b(\boldsymbol{\theta}, \boldsymbol{\theta}')} \\ &\leq \sqrt{s_b^2(\boldsymbol{\theta}) + s_b^2(\boldsymbol{\theta}') + 2s_b(\boldsymbol{\theta})s_b(\boldsymbol{\theta}')} \\ &= s_b(\boldsymbol{\theta}) + s_b(\boldsymbol{\theta}'). \end{aligned} \quad (\text{C.9})$$

We further bound (C.9) in terms of the  $b$  evaluations. We note that since the optimal method for minimising either the conditional or unconditional error also minimises  $\sigma_b(\boldsymbol{\theta}, \boldsymbol{\theta}')$ , any method for selecting the  $b$  evaluation locations will give an upper bound.

Suppose that  $\boldsymbol{\theta}_{1:t}$  consists of  $t \geq 1$  evaluation locations that all are  $\boldsymbol{\theta}$ . We denote  $\mathbf{1} \triangleq [1, \dots, 1]^\top \in \mathbb{R}^t$  and  $\mathbf{I}$  is a  $t \times t$  identity matrix. Then

$$\begin{aligned} s_t^2(\boldsymbol{\theta}) &= k(\boldsymbol{\theta}, \boldsymbol{\theta}) - k(\boldsymbol{\theta}, \boldsymbol{\theta}_{1:t})[k(\boldsymbol{\theta}_{1:t}, \boldsymbol{\theta}_{1:t}) + \sigma_n^2(\boldsymbol{\theta})\mathbf{I}]^{-1}k(\boldsymbol{\theta}_{1:t}, \boldsymbol{\theta}) \\ &= \sigma_s^2 - \sigma_s^2 \mathbf{1}^\top \left[ \mathbf{1}\mathbf{1}^\top + \frac{\sigma_n^2(\boldsymbol{\theta})}{\sigma_s^2} \mathbf{I} \right]^{-1} \mathbf{1}. \end{aligned} \quad (\text{C.10})$$

We use Sherman–Morrison formula to compute  $\left[ \mathbf{1}\mathbf{1}^\top + \frac{\sigma_n^2(\boldsymbol{\theta})}{\sigma_s^2} \mathbf{I} \right]^{-1} = \frac{\sigma_s^2}{\sigma_n^2(\boldsymbol{\theta})} \mathbf{I} - \frac{(\sigma_s^2/\sigma_n^2(\boldsymbol{\theta}))^2}{1 + t\sigma_s^2/\sigma_n^2(\boldsymbol{\theta})} \mathbf{1}\mathbf{1}^\top$ . After plugging this formula to (C.10) and some straightforward calculations, we see that

$$s_t(\boldsymbol{\theta}) = \frac{\sigma_s}{\sqrt{1 + t\sigma_s^2/\sigma_n^2(\boldsymbol{\theta})}} \leq \min \left\{ \sigma_s, \frac{\sigma_n(\boldsymbol{\theta})}{\sqrt{t}} \right\}, \quad (\text{C.11})$$

which also works when  $t = 0$  because then  $\sigma_n(\boldsymbol{\theta})/\sqrt{t} = \infty$  so that  $s_0(\boldsymbol{\theta}) = \sigma_s$ .

Suppose now that we have  $m \geq 0$  evaluations at  $\boldsymbol{\theta}$  and  $m' \geq 0$  evaluations at  $\boldsymbol{\theta}'$  such that  $m + m' \leq b$ . Then

$$\begin{aligned} s_b(\boldsymbol{\theta}) + s_b(\boldsymbol{\theta}') &\leq s_m(\boldsymbol{\theta}) + s_{m'}(\boldsymbol{\theta}') \\ &\stackrel{(\text{C.11})}{\leq} \min \left\{ \sigma_s, \frac{\sigma_n(\boldsymbol{\theta})}{\sqrt{m}} \right\} + \min \left\{ \sigma_s, \frac{\sigma_n(\boldsymbol{\theta}')}{\sqrt{m'}} \right\} \\ &\leq \min \left\{ 2\sigma_s, \left( \frac{1}{\sqrt{m}} + \frac{1}{\sqrt{m'}} \right) \bar{\sigma}_n \right\}, \end{aligned}$$

where the first inequality follows from the fact that the GP variance function is non-increasing as a set function of data, that is, adding more evaluation locations to data cannot increase the GP variance. If we choose in particular  $m = m' = \lfloor b/2 \rfloor$ , we obtain

$$s_b(\boldsymbol{\theta}) + s_b(\boldsymbol{\theta}') \leq 2 \min \left\{ \sigma_s, \bar{\sigma}_n / \sqrt{\lfloor b/2 \rfloor} \right\} = c_b. \quad (\text{C.12})$$



Combining (C.8), (C.9) and (C.12) produces the final bound (C.4).

To prove (C.5), we first notice that  $E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is decreasing function with respect to  $\mu_b$  when  $\mu_b \geq 0$  so that the choice  $\mu_b = 0$  maximises  $E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  in  $\mu_b \geq 0$ .

Suppose now  $\mu_b \leq 0$ . For this case we already derived the formula:

$$E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \Phi\left(\frac{\mu_b}{\sigma_b}\right) + e^{\mu_b + \sigma_b^2/2} \left( \Phi\left(-\frac{\mu_b + \sigma_b^2}{\sigma_b}\right) - 2\Phi(-\sigma_b) \right).$$

We maximise it with respect to  $\mu_b$  and use the inequality for  $\sigma_b$  derived in the first part of this proof and the fact that  $\sigma_b \mapsto E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is strictly increasing function for  $\sigma_b > 0$  (which we see directly from (16)) to obtain

$$\begin{aligned} E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') &\leq \max_{\mu \leq 0} \left\{ \Phi\left(\frac{\mu}{\sigma_b}\right) + e^{\mu + \sigma_b^2/2} \left( \Phi\left(-\frac{\mu + \sigma_b^2}{\sigma_b}\right) - 2\Phi(-\sigma_b) \right) \right\} \\ &\leq \max_{\mu \leq 0} \left\{ \Phi\left(\frac{\mu}{c_b}\right) + e^{\mu + c_b^2/2} \left( \Phi\left(-\frac{\mu + c_b^2}{c_b}\right) - 2\Phi(-c_b) \right) \right\}, \end{aligned}$$

which is the desired bound. ■

**Justification for (C.6).** We write

$$f(\boldsymbol{\theta}') - f(\boldsymbol{\theta}) = \log \mathcal{N}(\boldsymbol{\theta}' | \mathbf{0}, \boldsymbol{\Sigma}) - \log \mathcal{N}(\boldsymbol{\theta} | \mathbf{0}, \boldsymbol{\Sigma}) = \frac{1}{2} (\boldsymbol{\theta}'^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\theta}' - \boldsymbol{\theta}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\theta}). \quad (\text{C.13})$$

Since  $\boldsymbol{\Sigma}$  is positive definite, we have Cholesky factorisation  $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\top$  so that  $\boldsymbol{\Sigma}^{-1} = \mathbf{L}^{-\top} \mathbf{L}^{-1}$ , where  $\mathbf{L}^{-\top} \triangleq (\mathbf{L}^{-1})^\top = (\mathbf{L}^\top)^{-1}$ . Consider random vectors  $\boldsymbol{\psi} = \mathbf{L}^{-1} \boldsymbol{\theta}$  and  $\boldsymbol{\psi}' = \mathbf{L}^{-1} \boldsymbol{\theta}'$ . Clearly  $\boldsymbol{\theta}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\theta} = \boldsymbol{\psi}^\top \boldsymbol{\psi}$ ,  $\boldsymbol{\theta}'^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\theta}' = \boldsymbol{\psi}'^\top \boldsymbol{\psi}'$  and  $[\boldsymbol{\psi}^\top, \boldsymbol{\psi}'^\top]^\top$  is Gaussian distributed. We compute  $\mathbb{E}(\boldsymbol{\psi}) = \mathbf{L}^{-1} \mathbb{E}(\boldsymbol{\theta}) = \mathbf{0}$  and  $\mathbb{V}(\boldsymbol{\psi}) = \mathbf{L}^{-1} \mathbb{V}(\boldsymbol{\theta}) \mathbf{L}^{-\top} = \mathbf{L}^{-1} \boldsymbol{\Sigma} \mathbf{L}^{-\top} = \mathbf{I}$ . we also have  $\mathbb{E}(\boldsymbol{\psi}') = \mathbb{E}(\mathbb{E}(\mathbf{L}^{-1} \boldsymbol{\theta}' | \boldsymbol{\theta})) = \mathbb{E}(\mathbf{L}^{-1} \mathbb{E}(\boldsymbol{\theta}' | \boldsymbol{\theta})) = \mathbb{E}(\mathbf{L}^{-1} \boldsymbol{\theta}) = \mathbf{0}$  and

$$\begin{aligned} \mathbb{V}(\boldsymbol{\psi}') &= \mathbb{E}(\mathbb{V}(\mathbf{L}^{-1} \boldsymbol{\theta}' | \boldsymbol{\theta})) + \mathbb{V}(\mathbb{E}(\mathbf{L}^{-1} \boldsymbol{\theta}' | \boldsymbol{\theta})) \\ &= \mathbb{E}(\mathbf{L}^{-1} \mathbb{V}(\boldsymbol{\theta}' | \boldsymbol{\theta}) \mathbf{L}^{-\top}) + \mathbb{V}(\mathbf{L}^{-1} \mathbb{E}(\boldsymbol{\theta}' | \boldsymbol{\theta})) \\ &= \mathbb{E}(s^2 \mathbf{L}^{-1} \boldsymbol{\Sigma} \mathbf{L}^{-\top}) + \mathbb{V}(\mathbf{L}^{-1} \boldsymbol{\theta}) \\ &= s^2 \mathbf{I} + \mathbf{L}^{-1} \boldsymbol{\Sigma} \mathbf{L}^{-\top} \\ &= (s^2 + 1) \mathbf{I}. \end{aligned}$$

Since we can write  $\boldsymbol{\theta}' = \boldsymbol{\theta} + \mathbf{r}$ , where  $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, s^2 \boldsymbol{\Sigma})$ , it follows that

$$\begin{aligned} \text{cov}(\boldsymbol{\psi}, \boldsymbol{\psi}') &= \text{cov}(\mathbf{L}^{-1} \boldsymbol{\theta}, \mathbf{L}^{-1}(\boldsymbol{\theta} + \mathbf{r})) \\ &= \text{cov}(\mathbf{L}^{-1} \boldsymbol{\theta}, \mathbf{L}^{-1} \boldsymbol{\theta}) + \text{cov}(\mathbf{L}^{-1} \boldsymbol{\theta}, \mathbf{L}^{-1} \mathbf{r}) \\ &= \mathbb{V}(\mathbf{L}^{-1} \boldsymbol{\theta}) + \mathbf{L}^{-1} \text{cov}(\boldsymbol{\theta}, \mathbf{r}) \mathbf{L}^{-\top} \\ &= \mathbf{I}. \end{aligned}$$

We have thus shown

$$\begin{aligned} f(\boldsymbol{\theta}') - f(\boldsymbol{\theta}) &= \frac{1}{2} (\boldsymbol{\psi}^\top \boldsymbol{\psi} - \boldsymbol{\psi}'^\top \boldsymbol{\psi}') = \frac{1}{2} \sum_{i=1}^p (\psi_i^2 - \psi_i'^2), \\ \begin{bmatrix} \boldsymbol{\psi} \\ \boldsymbol{\psi}' \end{bmatrix} &\sim \mathcal{N}_{2p} \left( \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{I} & (s^2 + 1) \mathbf{I} \end{bmatrix} \right). \end{aligned} \quad (\text{C.14})$$

This shows that the distribution of  $f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})$  does not depend on  $\boldsymbol{\Sigma}$  and is approximately Gaussian by the central limit theorem (which applies because the random variables  $\psi_i^2 - \psi_i'^2, i = 1, \dots, p$  are independent and have finite variance by (C.14)) when  $p$  is large.

The expectation and variance are now obtained as<sup>4</sup>:

$$\begin{aligned}\mathbb{E}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) &= \frac{1}{2} \sum_{i=1}^p (\mathbb{E}(\psi_i^2) - \mathbb{E}(\psi_i'^2)) = -\frac{1}{2}ps^2, \\ \mathbb{V}_{\boldsymbol{\theta}, \boldsymbol{\theta}'}(f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})) &= \frac{1}{4} \sum_{i=1}^p (\mathbb{V}(\psi_i^2) + \mathbb{V}(\psi_i'^2) - 2\text{cov}(\psi_i^2, \psi_i'^2)) = \frac{1}{2}ps^2(s^2 + 2),\end{aligned}$$

where we have additionally used the facts  $\mathbb{V}(\psi_i^2) = \mathbb{E}(\psi_i^4) - \mathbb{E}(\psi_i^2)^2$ ,  $\text{cov}(\psi_i^2, \psi_i'^2) = \mathbb{E}(\psi_i^2\psi_i'^2) - \mathbb{E}(\psi_i^2)\mathbb{E}(\psi_i'^2)$  and well-known formulas for the moments of zero-mean Gaussian distribution.

**Derivation of the revised upper bounds.** Proposition C.1 shows the worst case upper bounds with respect to  $\mu_b$ . Here we derive revised bounds where we instead use a specific distribution for  $\mu_b$  under the Gaussian target and proposal assumption. That is, we assume  $\mu_b$  follows (for each possible  $b$ ) the same distribution as  $f(\boldsymbol{\theta}') - f(\boldsymbol{\theta})$  in (C.14). Under the assumptions of Proposition C.1 and using (C.7) we obtain

$$\begin{aligned}\mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \varepsilon) &= \int_{\mathbb{R}} \mathbb{P}(\mathcal{E}_{b,u,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \varepsilon \mid \mu_b) \pi(\mu_b) d\mu_b \\ &= \int_{\mathbb{R}} \left( \max\{1 - e^{\mu_b - \lambda_b}, 0\} + \min\{e^{\mu_b + \lambda_b} - 1, 0\} \right) \pi(\mu_b) d\mu_b \\ &\leq \int_{\mathbb{R}} \left( \max\{1 - e^{\mu_b + \Phi^{-1}(\varepsilon)c_b}, 0\} + \min\{e^{\mu_b - \Phi^{-1}(\varepsilon)c_b} - 1, 0\} \right) \pi(\mu_b) d\mu_b \\ &\approx \frac{1}{r} \sum_{i=1}^r \left( \max\{1 - e^{\mu_b^{(i)} + \Phi^{-1}(\varepsilon)c_b}, 0\} + \min\{e^{\mu_b^{(i)} - \Phi^{-1}(\varepsilon)c_b} - 1, 0\} \right),\end{aligned}\tag{C.15}$$

where  $\mu_b^{(i)} \sim_{\text{i.i.d.}} \pi(\mu_b), i = 1, \dots, r$  (for any  $b$ ). Samples from  $\pi(\mu_b)$  can be obtained by first drawing  $[\psi_j^{(i)}, \psi_j'^{(i)}]^\top \sim_{\text{i.i.d.}} \mathcal{N}_2 \left( \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 1 & s^2 + 1 \end{bmatrix} \right)$  for  $j = 1, \dots, p$  and then computing  $\mu_b^{(i)} = \sum_{j=1}^p (\psi_j^{(i)2} - \psi_j'^{(i)2})/2$ .

We write the unconditional error as  $E(\mu_b, \sigma_b)$  when it is considered as a function of  $\mu_b$  and  $\sigma_b$ . We then obtain

$$\mathbb{P}(E_{b,\hat{\gamma}}(\boldsymbol{\theta}, \boldsymbol{\theta}') \geq \varepsilon) = \mathbb{P}(E(\mu_b, \sigma_b) \geq \varepsilon) \leq \mathbb{P}(E(\mu_b, c_b) \geq \varepsilon) \approx \frac{1}{r} \sum_{i=1}^r \mathbf{1}_{E(\mu_b^{(i)}, c_b) \geq \varepsilon},\tag{C.16}$$

where  $c_b$  is as in Proposition C.1 and where  $\mu_b^{(i)} \sim_{\text{i.i.d.}} \pi(\mu_b), i = 1, \dots, r$  can be simulated as above.

---

4. One could also write  $\boldsymbol{\psi}^\top \boldsymbol{\psi} - \boldsymbol{\psi}'^\top \boldsymbol{\psi}' = [\boldsymbol{\psi}^\top, \boldsymbol{\psi}'^\top] \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\psi} \\ \boldsymbol{\psi}' \end{bmatrix}$  and then use known formulas for computing the expectation and variance of this quadratic form to obtain the same results.

## Appendix D. Additional illustrations

Figures D.1 and D.2 show how a GP prior with a non-zero mean function and an additional evaluation near the right boundary of the parameter space, respectively, produce more intuitive estimates of the SL posterior in the illustrative example of Section 4.2.2. In the former case large uncertainty of the likelihood near the right boundary however remains.




Figure D.1: As in Figure 2 but here a GP prior with a special mean function  $m_0(\theta) = \beta_1 + \beta_2\theta + \beta_3\theta^2$  as described in Section 2.5 is used instead of a zero-mean GP prior.




Figure D.2: As in Figure 2 but here an additional evaluation at  $\theta = 25$  is used for GP fitting.

## Appendix E. Additional discussion on modelling assumptions and implementation

We make some remarks on the Gaussian noise assumption (9) in Section E.1. In Section E.2 we then discuss how potential practical difficulties with GP fitting are handled in our GP-MH algorithm.

### E.1 Remarks on the Gaussian noise assumption

An unbiased, Gaussian distributed estimate of the *log*-likelihood function is ideal for GP-MH. Such approximately Gaussian estimate is available in the likelihood-free generalised Bayesian updating case of Section 5.3. The SL estimate of Section A.1 is in general not unbiased even if the Gaussianity assumption holds. The numerical results in Järvenpää

et al. (2021) however suggest that the resulting approximation error is small for typical  $N$  and that the Gaussian noise assumption is sensible.

An unbiased estimator of the (approximate) likelihood is available in some applications such as in the standard ABC case of Section A.2. Also, an unbiased estimate for the likelihood of the parameters of a state-space model with intractable transition density can be obtained via bootstrap particle filtering. The ABC estimate is considered by Wilkinson (2014) and the state-space model case by Drovandi et al. (2018) in the context of GP surrogate modelling but without clear justifications.

Suppose  $\mathbb{E}(p(\boldsymbol{\theta})) = \pi_{\text{lik}}(\boldsymbol{\theta}) \triangleq \pi(\mathbf{x}_o | \boldsymbol{\theta})$  where the expectation is wrt. the noise in the evaluation. Let  $y(\boldsymbol{\theta}) \triangleq \log p(\boldsymbol{\theta})$ . In general, we have

$$\mathbb{E}(y(\boldsymbol{\theta})) = \mathbb{E}(\log p(\boldsymbol{\theta})) \neq \log \mathbb{E}(p(\boldsymbol{\theta})) = \log \pi_{\text{lik}}(\boldsymbol{\theta}),$$

that is, an unbiased estimator of the likelihood function cannot be transformed to an unbiased estimate of the log-likelihood by just taking the logarithm. Similar observations, but in the opposite direction and in the context of pseudo-marginal MH, are made in Bardenet et al. (2017, Section 4.2) and Llorente et al. (2021, Section 2).

Suppose further that  $p(\boldsymbol{\theta})$  follows log-Normal distribution with the mean  $\pi_{\text{lik}}(\boldsymbol{\theta})$  and variance  $\sigma^2(\boldsymbol{\theta}) \geq 0$ . Then we have

$$y(\boldsymbol{\theta}) \sim \mathcal{N}\left(\log\left(\frac{\pi_{\text{lik}}(\boldsymbol{\theta})^2}{\sqrt{\pi_{\text{lik}}(\boldsymbol{\theta})^2 + \sigma^2(\boldsymbol{\theta})}}\right), \log\left(1 + \frac{\sigma^2(\boldsymbol{\theta})}{\pi_{\text{lik}}(\boldsymbol{\theta})^2}\right)\right)$$

by the basic properties of the log-Normal distribution. It follows that

$$\mathbb{E}(y(\boldsymbol{\theta})) = \log \pi_{\text{lik}}(\boldsymbol{\theta}) + \log\left(\frac{\pi_{\text{lik}}(\boldsymbol{\theta})}{\sqrt{\pi_{\text{lik}}(\boldsymbol{\theta})^2 + \sigma^2(\boldsymbol{\theta})}}\right) \leq \log \pi_{\text{lik}}(\boldsymbol{\theta}).$$

Hence, if  $\sigma^2(\boldsymbol{\theta}) \ll \pi_{\text{lik}}(\boldsymbol{\theta})^2$  in this case, then  $y(\boldsymbol{\theta})$  might work as a reasonable approximation for an unbiased estimate of  $\log \pi_{\text{lik}}(\boldsymbol{\theta})$ .

## E.2 Implementation details on handling log-likelihood evaluations

As mentioned in the main text, the log-likelihood function can behave irregularly in some boundary regions of the parameter space. This is usually not problematic for standard MH (unless one tries to initialise MH from such a region) because any proposal that results an infeasible log-likelihood evaluation is simply rejected. However, handling such log-likelihood values in GP-based methods requires care because including individual values with substantially different magnitudes to  $\mathcal{D}_t$  typically leads to poor GP fits or numerical issues. In B(O)LFI methods one would additionally need to somehow ensure that the global optimum of the acquisition function does not lie on such problematic parameter regions where the GP fit cannot be trusted. In the following we discuss how this practical problem is handled in our GP-MH implementation.

We treat a log-likelihood evaluation  $y_j$  at  $\boldsymbol{\theta}_j \in \Theta$  as *invalid* in our algorithms if any of the following holds:  $y_j$  is a complex number or “NaN” (not-a-number),  $|y_j| > 10^5$  or  $\sigma_n(\boldsymbol{\theta}_j) > 10^3$ . Invalid  $(y_j, \boldsymbol{\theta}_j)$  is never included to  $\mathcal{D}_t$  and hence not used for GP fitting.

Invalid evaluations can result at different stages of GP-MH. First, if  $2t_{\text{init}}$  tries do not produce  $t_{\text{init}}$  valid initial evaluations, the algorithm is terminated as of having too poor initialisation. If an invalid evaluation is obtained at the proposed point, the proposal is rejected and the algorithm continues as normal (neither  $\mathcal{Z}_t$  nor GP is updated). If the invalid evaluation is obtained at the current point then this means that the algorithm has likely proceeded to a point which should ideally have been rejected in an earlier iteration but were not based on the GP model. Because it may take long before the algorithm would manage to exit the problematic region, in this case the algorithm is terminated. EPoE often evaluates at a point which is neither the current nor the proposed point. If an invalid evaluation is encountered in this case, it is neglected and a new evaluation is obtained using naive instead so that the algorithm can proceed as in the naive/EPoE case. Unless terminated, the algorithm switches back to EPoE after this exceptional step.

The above heuristic procedure often allows the algorithm to recover if a problematic likelihood evaluation is obtained. A user can rerun the algorithm using a better initialisation if it was terminated. Note that technically any parameter can produce an invalid evaluation under the Gaussian noise assumption (9). Apart from some unusual or pathological situations, this is however expected to happen so rarely in the highest density region that the potential bias caused by always neglecting a proposal that produced an invalid evaluation is not taken into account. A practical difficulty is that the range of the log-likelihood function in the highest density region is rarely known in advance and depends, among other things, on the scaling of the data space, possible model misspecification and whether potential additive constant terms in the log-likelihood formula (which cancel out in the MH accept/reject test but still affect the GP modelling) are neglected. Our definition of the invalid log-likelihood evaluation is based on our numerical experiments and some simple analyses but may require adjustments in some other settings due to these reasons.

## Appendix F. Additional details on noisy synthetic densities

We here summarise the details of the three 6D toy log-densities originally presented in Järvenpää et al. (2021) (where also their 2D versions were used for illustration and are shown as Figure D.2 of their supplementary material) and used in Section 5.1 of this article. These log-densities, which we denote as  $f_{6\text{D}}$ , are constructed so that  $f_{6\text{D}}(\boldsymbol{\theta}) = f_{2\text{D}}(\boldsymbol{\theta}_{1:2}) + f_{2\text{D}}(\boldsymbol{\theta}_{3:4}) + f_{2\text{D}}(\boldsymbol{\theta}_{5:6})$ . The 2D log densities  $f_{2\text{D}}$  are then defined so that the ‘Simple’ log-density results when  $f_{2\text{D}}(\boldsymbol{\theta}) = -\boldsymbol{\theta}^\top \mathbf{S}_\rho^{-1} \boldsymbol{\theta} / 2$  where  $\rho = 0.25$ , the ‘Banana’ results when  $f_{2\text{D}}(\boldsymbol{\theta}) = -[\theta_1, \theta_2 + \theta_1^2 + 1] \mathbf{S}_\rho^{-1} [\theta_1, \theta_2 + \theta_1^2 + 1]^\top / 2$  where  $\rho = 0.9$  and, finally, the ‘Multimodal’ log-density is obtained using  $f_{2\text{D}}(\boldsymbol{\theta}) = -[\theta_1, \theta_2^2 - 2] \mathbf{S}_\rho^{-1} [\theta_1, \theta_2^2 - 2]^\top / 2$  where  $\rho = 0.5$ . Above we have defined  $\mathbf{S}_\rho \in \mathbb{R}^{2 \times 2}$  so that  $(S_\rho)_{11} = (S_\rho)_{22} = 1$  and  $(S_\rho)_{12} = (S_\rho)_{21} = \rho$ . The 2D structure of these models is used to aid computing the ground-truth posterior but is not taken into account in the GP modelling.

The above log-densities are additionally modified by specifying bounds for their six parameters. In practice this is achieved by using the above log-densities in the place of the target log-likelihood function for the GP-MH algorithm and by setting the uniform priors  $\mathcal{U}([-16, 16]^6)$ ,  $\mathcal{U}(\prod_{i=1}^3 ([-6, 6] \times [-20, 2]))$  and  $\mathcal{U}([-6, 6]^6)$  for Simple, Banana and Multimodal densities, respectively. We use the following initial points in our experiments:

$\theta^{(0)} = -8\mathbf{1}$  for Simple and  $\theta^{(0)} = -3\mathbf{1}$  for both Banana and Multimodal. The initial covariance matrix of the Gaussian proposal is  $\Sigma_0 = \mathbf{I}$  for all three test cases.

## Appendix G. Additional results and experiments

### G.1 Noisy synthetic densities

We here complement Section 5.1 with additional results. First, Figure G.1 demonstrates how the quality of the marginal posterior approximation of GP-MH and MH-BLFI develops as a function of iteration  $i$ , that is, as more approximate MH samples are obtained. We observe that  $10^5$  iterations are enough for Banana and Multimodal while  $10^4$  iterations is already sufficient for the Simple log-density. Since the results by MH-BLFI, shown in the bottom row of Figure G.1, are based on a separate MH sampling with chain length  $10^5$ , its convergence is not directly affected by  $i$ . The MH-BLFI results during the initial iterations are nevertheless poor ( $\text{TV} \approx 1$ ) because the number of collected log-likelihood evaluations is obviously small initially and the resulting GP approximation hence inaccurate. Both methods eventually produce approximations with similar quality with each  $\varepsilon$ .

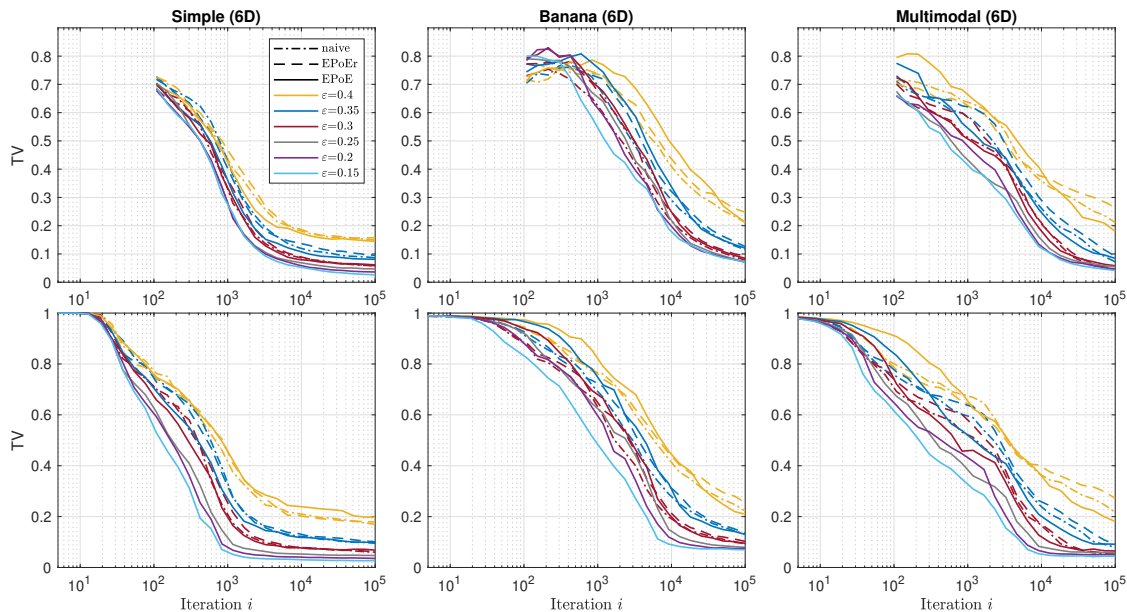


Figure G.1: Accuracy of the marginal posterior approximation as a function of iteration  $i$  of Algorithm 3. Each line shows the median result over the 100 repeats. *Top row* shows GP-MH and the *bottom row* the corresponding results by MH-BLFI. The results by GP-MH at the early iterations ( $i < 10^2$ ) are not shown because sampling error is necessarily large there.

Figure G.2 shows that most of the log-density evaluations of the experiments of Section 5.1 occur already in the early stage of the algorithm which is neglected as burn-in. This is also the case with our other experiments and explains why GP-MH and MH-BLFI produce similar results. The EPoE results with  $\varepsilon = 0.25, 0.2$  and  $0.15$  are not shown as they were very similar to those with  $\varepsilon = 0.3$ .




Figure G.2: Proportion of the total log-density evaluations collected as a function of GP-MH iteration  $i$ . *Top row*: Simple (6D), *middle row*: Banana (6D), *bottom row*: Multimodal (6D) test density. The solid lines show the median and the dashed lines the 75% quantile computed over the 100 repeated runs.

Figure G.3 demonstrates the quality of the marginal posterior approximation as in Section 5.1 but when the noise levels have been doubled to  $\sigma_n = 4$  for Simple and  $\sigma_n = 2$  for Banana and Multimodal. All methods still produce reasonable results but more evaluations are of course needed. Our threshold for the maximum number of evaluations  $10^3$  is always met in Banana and Multimodal cases using EPoEr and naive with  $\varepsilon = 0.3$  or EPoE with  $\varepsilon = 0.15$ . EPoE is clearly more sample-efficient than EPoEr and naive. EPoE produces similar or slightly better sample-efficiency than BLFI with IMIQR.




Figure G.3: Accuracy of the marginal posterior approximation as a function of the required log-density evaluations at the final iteration  $i_{\text{MH}} = 10^5$ . These results are as in Figure 3 except that the noise levels were doubled and only the GP-MH results are shown.

## G.2 Ricker model

We briefly consider the (scaled) Ricker model used before e.g. by Wood (2010); Gutmann and Corander (2016); Price et al. (2018) to analyse LFI methods. In this model the population evolves as

$$N_{t+1} = rN_t \exp(-N_t + \varepsilon_t),$$

for  $t = 1, \dots, T$ . The process noise and observation models are as in Section 5.2.1. Our experimental set-up is the same as for the theta-Ricker model except for the following differences: The priors are  $\log(r) \sim \mathcal{U}([3, 5])$ ,  $\phi \sim \mathcal{U}([4, 20])$ ,  $\sigma_\varepsilon \sim \mathcal{U}([0, 0.8])$  and the “true” parameter is  $\theta_{\text{true}} = (3.8, 10, 0.3)$ . We set  $T = 50$ ,  $t_{\text{init}} = 10$ ,  $i_{\text{MH}} = 10^5$ ,  $\theta^{(0)} = (3.4, 8.0, 0.15)$  and  $\Sigma_0 = \text{diag}(0.1, 1.0, 0.1)^2$ .

Figure G.4 shows the results in a similar fashion as before. We observe that less log-SL evaluations are needed as for theta-Ricker. BLFI with IMIQR produces slightly better sample-efficiency and accuracy than the GP-MH methods in this experiment. Figure G.5 shows a typical example of the estimated posterior obtained using  $\varepsilon = 0.2$  and EPoE. We can see that both the marginals and the correlation structure is estimated well.

## G.3 Theta-Ricker model

We complement the theta-Ricker experiments of Section 5.2.1 with some additional results. Figure G.6 demonstrates that naive, EPoEr and EPoE all tend to produce fairly similar evaluation locations. The evaluations by EPoE are however slightly more evenly distributed and diverse as those by the other two methods.



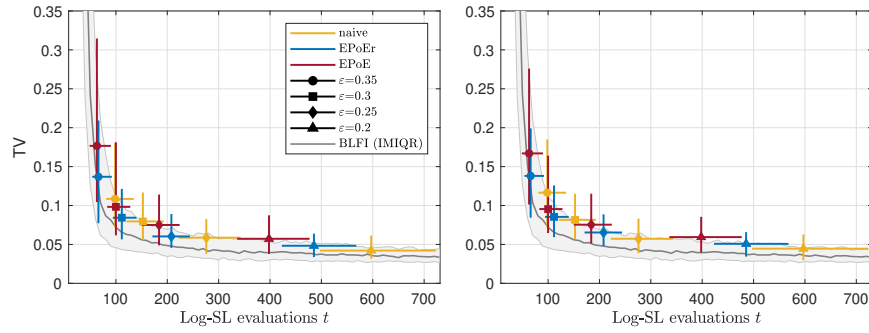


Figure G.4: Marginal posterior approximation accuracy as in Figure 5 but for Ricker experiment and at  $i_{MH} = 10^5$  iterations. *Left plot*: GP-MH, *right plot*: MH-BLFI.

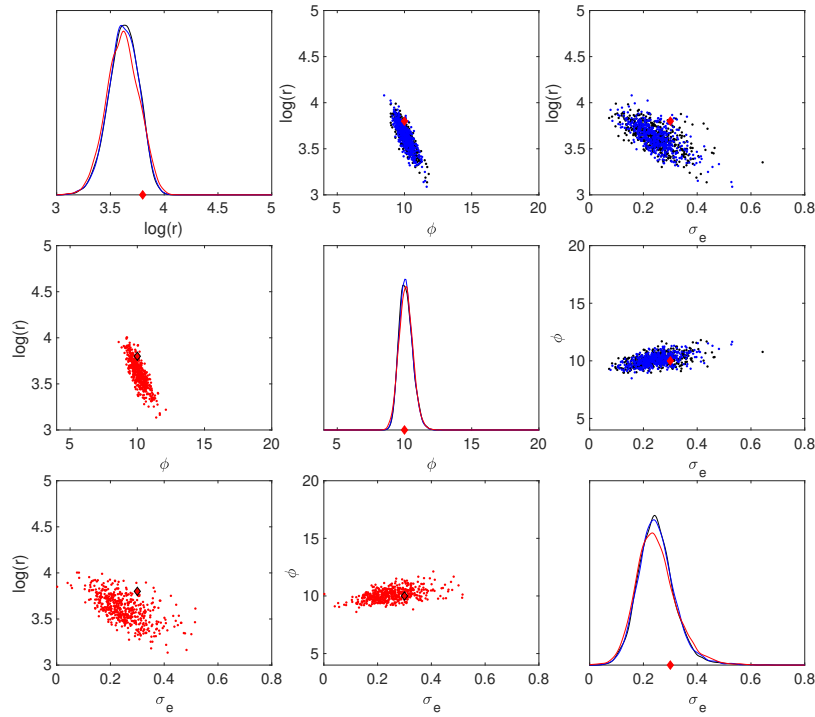


Figure G.5: Comparison of the ground-truth posterior (red dots/line) and a typical example of estimated posterior (black and blue dots/line) in the case of Ricker experiment. See the caption of Figure 6 for more detailed description.

Figure G.7 shows how the quality of the posterior approximation develops as a function of iteration  $i$  of Algorithm 3. Finally, Figure G.8 demonstrates, similarly to Figure G.2, the proportion of the collected log-SL evaluations.

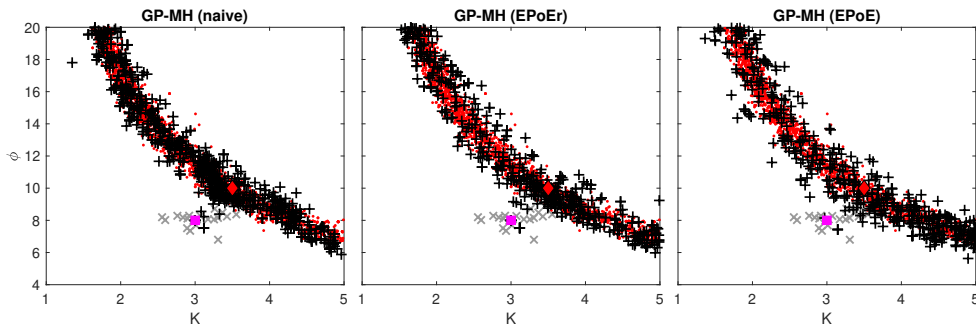


Figure G.6: Typical realisations of the log-SL evaluation locations projected to  $(K, \phi)$ -space in the theta-Ricker experiment. See the caption of Figure 4 for other details.

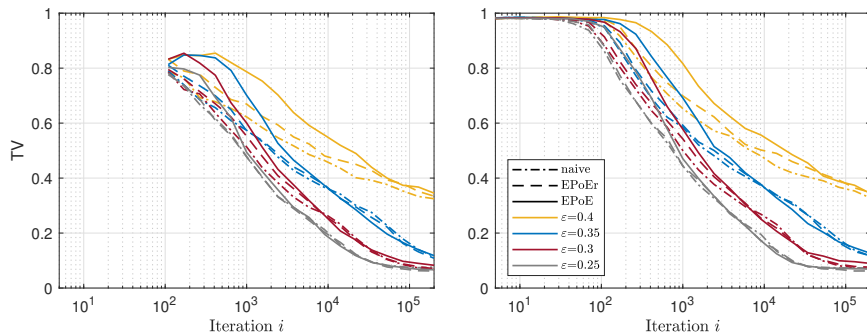


Figure G.7: Accuracy of the marginal posterior approximation in the theta-Ricker experiment as a function of iteration  $i$  of Algorithm 3. *Left plot*: GP-MH, *right plot*: MH-BLFI.




Figure G.8: Proportion of the total log-SL evaluations collected as a function of GP-MH iteration  $i$  in the theta-Ricker experiment. The solid lines show the median and the dashed lines the 75% quantile computed over the 100 repeated runs.

## References

- L. Acerbi. Variational Bayesian Monte Carlo. In *Advances in Neural Information Processing Systems 31*, pages 8223–8233, 2018.
- L. Acerbi. Variational Bayesian Monte Carlo with Noisy Likelihoods. In *Advances in Neural Information Processing Systems 33*, 2020.
- L. Alawieh, J. Goodman, and J. B. Bell. Iterative construction of Gaussian process surrogate models for Bayesian inference. *Journal of Statistical Planning and Inference*, 207:55–72, 2020.
- Z. An, L. F. South, D. J. Nott, and C. C. Drovandi. Accelerating Bayesian synthetic likelihood with the graphical lasso. *Journal of Computational and Graphical Statistics*, 28(2):471–475, 2019.
- Z. An, D. J. Nott, and C. C. Drovandi. Robust Bayesian synthetic likelihood via a semi-parametric approach. *Statistics and Computing*, 30:543–557, 2020.
- E. Angelino, M. J. Johnson, and R. P. Adams. Patterns of Scalable Bayesian Inference. *Foundations and Trends in Machine Learning*, 9(2-3):119–247, 2016.
- R. Bardenet, A. Doucet, and C. Holmes. On Markov chain Monte Carlo methods for tall data. *Journal of Machine Learning Research*, 18(47):1–43, 2017.
- J. Bect, D. Ginsbourger, L. Li, V. Picheny, and E. Vazquez. Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing*, 22(3):773–793, 2012.
- P. G. Bissiri, C. C. Holmes, and S. G. Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130, 2016.
- N. Bliznyuk, D. Ruppert, C. Shoemaker, R. Regis, S. Wild, and P. Mugunthan. Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation. *Journal of Computational and Graphical Statistics*, 17(2):270–294, 2008.
- F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic. Probabilistic integration: A role in statistical computation? *Statistical Science*, 34(1):1–22, 2019.
- H. R. Chai and R. Garnett. Improving quadrature for constrained integrands. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2751–2759, 2019.
- J. A. Christen and C. Fox. Markov Chain Monte Carlo Using an Approximation. *Journal of Computational and Graphical Statistics*, 14(4):795–810, 2005.
- E. Cleary, A. Garbuno-Inigo, S. Lan, T. Schneider, and A. M. Stuart. Calibrate, emulate, sample. *Journal of Computational Physics*, 424, 2021.

- J. Cockayne, C. Oates, T. Sullivan, and M. Girolami. Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789, 2019.
- P. R. Conrad, Y. M. Marzouk, N. S. Pillai, and A. Smith. Accelerating Asymptotically Exact MCMC for Computationally Intensive Models via Local Approximations. *Journal of the American Statistical Association*, 111(516):1591–1607, 2016.
- K. Cranmer, J. Brehmer, and G. Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020.
- C. C. Drovandi, M. T. Moores, and R. J. Boys. Accelerating pseudo-marginal MCMC using Gaussian processes. *Computational Statistics & Data Analysis*, 118:1–17, 2018.
- M. Fasiolo, N. Pya, and S. N. Wood. A Comparison of Inferential Methods for Highly Nonlinear State Space Models in Ecology and Epidemiology. *Statistical Science*, 31(1):96–118, 2016.
- M. Fielding, D. J. Nott, and S-Y. Liang. Efficient MCMC Schemes for Computationally Expensive Posterior Distributions. *Technometrics*, 53(1):16–28, 2011.
- D. T. Frazier, D. J. Nott, C. Drovandi, and R. Kohn. Bayesian inference using synthetic likelihood: Asymptotics and adjustments. *Journal of the American Statistical Association*, 0(0):1–12, 2022.
- P. I. Frazier. A Tutorial on Bayesian Optimization, 2018. Available at <https://arxiv.org/abs/1807.02811>.
- R. Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.
- A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian data analysis*. Chapman & Hall/CRC Texts in Statistical Science, third edition, 2013.
- T. Gunter, M. A. Osborne, R. Garnett, P. Hennig, and S. J. Roberts. Sampling for inference in probabilistic models with fast Bayesian quadrature. In *Advances in Neural Information Processing Systems 27*, pages 2789–2797, 2014.
- M. U. Gutmann and J. Corander. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(125):1–47, 2016.
- H. Haario, E. Saksman, and J. Tamminen. An adaptive Metropolis algorithm. *Bernoulli*, 7(2):223–242, 2001.
- W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):97–109, 1970.
- P. Hennig and C. J. Schuler. Entropy Search for Information-Efficient Global Optimization. *Journal of Machine Learning Research*, 13(1999):1809–1837, 2012.

- P. Hennig, M. A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 471(2179):20150142, 2015.
- P. Hennig, M. A. Osborne, and H. P. Kersting. *Probabilistic Numerics: Computation as Machine Learning*. Cambridge University Press, 2022.
- M. Järvenpää, M. U. Gutmann, A. Pleska, A. Vehtari, and P. Marttinen. Efficient acquisition rules for model-based approximate Bayesian computation. *Bayesian Analysis*, 14(2):595–622, 2019.
- M. Järvenpää, A. Vehtari, and P. Marttinen. Batch simulations and uncertainty quantification in Gaussian process surrogate approximate Bayesian computation. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 779–788, 2020.
- M. Järvenpää, M. U. Gutmann, A. Vehtari, and P. Marttinen. Parallel Gaussian process surrogate Bayesian inference with noisy likelihood evaluations. *Bayesian Analysis*, 16(1):147–178, 2021.
- K. Kandasamy, J. Schneider, and B. Póczos. Query efficient posterior estimation in scientific experiments via Bayesian active learning. *Artificial Intelligence*, 243:45–56, 2017.
- A. Korattikara, Y. Chen, and M. Welling. Austerity in MCMC Land: Cutting the Metropolis-Hastings Budget. In *Proceedings of the 31st International Conference on Machine Learning*, pages 181–189, 2014.
- S. Lan, T. Bui-Thanh, M. Christie, and M. Girolami. Emulation of higher-order tensors in manifold Monte Carlo methods for Bayesian Inverse Problems. *Journal of Computational Physics*, 308:81–101, 2016.
- F. Llorente, L. Martino, J. Read, and D. Delgado. A survey of Monte Carlo methods for noisy and costly densities with application to reinforcement learning, 2021. Available at <https://arxiv.org/abs/2108.00490>.
- X. Lyu, M. Binois, and M. Ludkovski. Evaluating Gaussian process metamodels and sequential designs for noisy level set estimation. *Statistics and computing*, 31(43), 2021.
- J. M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the United States of America*, 100(26):15324–8, 2003.
- E. Meeds and M. Welling. GPS-ABC: Gaussian Process Surrogate Approximate Bayesian Computation. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, 2014.

- E. Numminen, L. Cheng, M. Gyllenberg, and J. Corander. Estimating the transmission dynamics of streptococcus pneumoniae from strain prevalence data. *Biometrics*, 69(3): 748–757, 2013.
- A. O’Hagan. Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 1991.
- A. O’Hagan and J. F. C. Kingman. Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40(1):1–42, 1978.
- M. A. Osborne, D. Duvenaud, R. Garnett, C. E. Rasmussen, S. J. Roberts, and Z. Ghahramani. Active Learning of Model Evidence Using Bayesian Quadrature. *Advances in Neural Information Processing Systems 26*, pages 1–9, 2012.
- D. B. Owen. Tables for computing bivariate normal probabilities. *The Annals of Mathematical Statistics*, 27(4):1075–1090, 12 1956.
- D. B. Owen. A table of normal integrals. *Communications in Statistics - Simulation and Computation*, 9(4):389–419, 1980.
- L. Pacchiardi and R. Dutta. Generalized Bayesian Likelihood-Free Inference Using Scoring Rules Estimators, 2021. Available at <https://arxiv.org/abs/2104.03889>.
- L. M. Paun and D. Husmeier. Emulation-Accelerated Hamiltonian Monte Carlo Algorithms for Parameter Estimation and Uncertainty Quantification in Differential Equation Models. *Statistics and Computing*, 32(1), 2022.
- U. Picchini, U. Simola, and J. Corander. Sequentially Guided MCMC Proposals for Synthetic Likelihoods and Correlated Synthetic Likelihoods. *Bayesian Analysis*, pages 1–31, 2022.
- L. Polansky, P. de Valpine, J. O. Lloyd-Smith, and W. M. Getz. Likelihood ridges and multimodality in population growth rate models. *Ecology*, 90(8):2313–2320, 2009.
- L. F. Price, C. C. Drovandi, A. Lee, and D. J. Nott. Bayesian synthetic likelihood. *Journal of Computational and Graphical Statistics*, 27(1):1–11, 2018.
- C. E. Rasmussen. Gaussian Processes to Speed up Hybrid Monte Carlo for Expensive Bayesian Integrals. *Bayesian Statistics 7*, pages 651–659, 2003.
- C. E. Rasmussen and Z. Ghahramani. Bayesian Monte Carlo. In *Advances in Neural Information Processing Systems 15*, pages 505–512, 2003.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- C. P Robert. *The Bayesian Choice*. Springer, New York, second edition, 2007.
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, New York, second edition, 2004.

- C. Sherlock, A. Golightly, and D. A. Henderson. Adaptive, Delayed-Acceptance MCMC for Targets With Expensive Likelihoods. *Journal of Computational and Graphical Statistics*, 26(2):434–444, 2017.
- S. Sisson, Y. Fan, and M. Beaumont. *Handbook of Approximate Bayesian Computation*. New York: Chapman and Hall/CRC, 2019.
- O. Thomas, R. Dutta, J. Corander, S. Kaski, and M. U. Gutmann. Likelihood-Free Inference by Ratio Estimation. *Bayesian Analysis*, 17(1):1–31, 2022.
- J. Vanhatalo, J. Riihimäki, J. Hartikainen, P. Jylänki, V. Tolvanen, and A. Vehtari. GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14: 1175–1179, 2013.
- H. Wang and J. Li. Adaptive Gaussian process approximation for Bayesian inference with expensive likelihood functions. *Neural Computation*, 30(11):3072–3094, 2018.
- R. D. Wilkinson. Accelerating ABC methods using Gaussian processes. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, 2014.
- J. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *Proceedings of the 37th International Conference on Machine Learning*, pages 10292–10302, 2020.
- S. Wqvist, U. Picchini, J. L. Forman, K. Lindorff-Larsen, and W. Boomsma. Accelerating delayed-acceptance Markov chain Monte Carlo algorithms, 2018. Available at <https://arxiv.org/abs/1806.05982>.
- S. N. Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466:1102–1104, 2010.
- C. Zhang, B. Shahbaba, and H. Zhao. Hamiltonian Monte Carlo acceleration using surrogate functions with random bases. *Statistics and Computing*, 27:1473–1490, 2017.
- J. Zhang and A. A. Taflanidis. Accelerating MCMC via Kriging-based adaptive independent proposals and delayed rejection. *Computer Methods in Applied Mechanics and Engineering*, 355:1124–1147, 2019.
- R. Zhang, A. F. Cooper, and C. M. De Sa. Asymptotically optimal exact minibatch Metropolis-Hastings. In *Advances in Neural Information Processing Systems*, volume 33, pages 19500–19510, 2020.