

Parallel Gaussian Process Surrogate Bayesian Inference with Noisy Likelihood Evaluations

Marko Järvenpää*, Michael U. Gutmann[†], Aki Vehtari* and Pekka Marttinen*

Abstract. We consider Bayesian inference when only a limited number of noisy log-likelihood evaluations can be obtained. This occurs for example when complex simulator-based statistical models are fitted to data, and synthetic likelihood (SL) method is used to form the noisy log-likelihood estimates using computationally costly forward simulations. We frame the inference task as a sequential Bayesian experimental design problem, where the log-likelihood function is modelled with a hierarchical Gaussian process (GP) surrogate model, which is used to efficiently select additional log-likelihood evaluation locations. Motivated by recent progress in the related problem of batch Bayesian optimisation, we develop various batch-sequential design strategies which allow to run some of the potentially costly simulations in parallel. We analyse the properties of the resulting method theoretically and empirically. Experiments with several toy problems and simulation models suggest that our method is robust, highly parallelisable, and sample-efficient.

Keywords: expensive likelihoods, likelihood-free inference, surrogate modelling, Gaussian processes, sequential experiment design, parallel computing.

1 Introduction

When the analytic form of the likelihood function of a statistical model is available, standard sampling techniques such as Markov Chain Monte Carlo (MCMC, see e.g. Robert and Casella 2004) can often be used for Bayesian inference. However, many models of interest in several areas of science, for example in computational biology and ecology, have an expensive-to-evaluate or intractable likelihood function which severely complicates inference. When the likelihood is intractable but forward simulation of the model is feasible, simulation-based inference methods (also called likelihood-free inference) such as approximate Bayesian computation (ABC) can be used. Unfortunately, such algorithms typically require a huge number of simulations making inference computationally costly. Examples of models with intractable likelihoods can be found in e.g. Beaumont et al. (2002); Marin et al. (2012); Lintusaari et al. (2017); Marttinen et al. (2015); Järvenpää et al. (2018) and Section 6.2 of this article.

Surrogate models, also called meta-models or emulators, such as Gaussian processes (Rasmussen and Williams, 2006) have been used extensively to calibrate deterministic computer codes, see e.g. Kennedy and O’Hagan (2001). GP surrogates have recently also

Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University, marko.jarvenpaa@aalto.fi, aki.vehtari@aalto.fi, pekka.marttinen@aalto.fi

[†]School of Informatics, University of Edinburgh, michael.gutmann@ed.ac.uk

been used to accelerate Bayesian inference by modelling some part of the inferential process, such as the log-likelihood function. The model allows extracting information from the simulations efficiently, and can be used e.g. to determine where additional simulations are needed. For example, Rasmussen (2003); Kandasamy et al. (2015); Sinsbeck and Nowak (2017); Drovandi et al. (2018); Wang and Li (2018); Acerbi (2018) have developed GP-based techniques to accelerate Bayesian inference when the exact likelihood or the corresponding deterministic model is tractable but expensive. Various GP surrogate techniques have been proposed also for ABC, where one can only draw samples i.e. pseudo-data from a statistical model but not evaluate the likelihood. These include Meeds and Welling (2014); Jabot et al. (2014); Wilkinson (2014); Gutmann and Corander (2016); Järvenpää et al. (2019).

In this paper we focus on GP surrogate modelling of noisy log-likelihood evaluations. Earlier works on emulating the log-likelihood function have mostly assumed exact, i.e., noiseless evaluations or the noise has not been explicitly modelled. We show that noisy evaluations cause extra challenges. Also, although not the focus of this work, we remark that one often has some control over the noise level. While our approach is applicable whenever noisy, expensive log-likelihood evaluations of a statistical model of interest are available, we mainly focus on likelihood-free inference using the synthetic likelihood method (Wood, 2010; Price et al., 2018), where the intractable log-likelihood is approximated using repeated forward simulations at each evaluation location.

Recently, Järvenpää et al. (2019) developed a Bayesian decision theoretic framework for ABC inference and considered sequential strategies (also called active learning) to select the next evaluation location for an expensive simulation model. Here we extend this framework in two ways: 1) we modify it to address the problem of Bayesian inference using noisy log-likelihood evaluations, which is different from ABC, and 2) we develop batch-sequential design strategies to efficiently parallelise the estimation of the surrogate likelihood. In earlier related works the simulation locations have been selected either sequentially (Kandasamy et al., 2015; Sinsbeck and Nowak, 2017; Wang and Li, 2018; Acerbi, 2018; Järvenpää et al., 2019) or using simple heuristics (Wilkinson, 2014; Gutmann and Corander, 2016). Batch strategies are useful when a computing cluster is available and, as we show, can substantially reduce the computation time compared to the corresponding sequential strategies. We also analyse some properties of the proposed methods theoretically, and conduct an extensive empirical comparison.

Our approach is closely related to Bayesian quadrature (BQ), see e.g. O’Hagan (1991); Hennig et al. (2015); Karvonen et al. (2018). In particular, BQ methods have been used by Osborne et al. (2012); Gunter et al. (2014); Chai and Garnett (2019) to compute the marginal likelihood and to quantify the numerical error of this integral probabilistically. In this article we are not interested in this particular integral but in obtaining an accurate estimate of the posterior. Also, we allow noisy log-likelihood evaluations. Another related problem is Bayesian optimisation (BO), see e.g. Brochu et al. (2010); Shahriari et al. (2015). Our objective to parallelise simulations is motivated by recent research on batch Bayesian optimisation (Ginsbourger et al., 2010; Azimi et al., 2010; Snoek et al., 2012; Contal et al., 2013; Desautels et al., 2014; Shah and Ghahramani, 2015; Wu and Frazier, 2016; Gonzalez et al., 2016; Wilson et al., 2018).

However, while BO methods can be used to accelerate likelihood-free Bayesian inference (Gutmann and Corander, 2016), they are not specifically designed for estimating the posterior (see discussion in e.g. Kandasamy et al. 2015; Järvenpää et al. 2019). Similarly to Järvenpää et al. (2019), we explicitly design our algorithms from the first principles of Bayesian decision theory to acknowledge the goal of the analysis, i.e. estimation of the posterior density. Finally, we note that GPs in conjunction with Bayesian experimental designs have also been successful in estimating level and excursion sets of expensive-to-evaluate functions, see e.g. Bect et al. (2012); Chevalier et al. (2014); Lyu et al. (2018).

This paper is organised as follows. Section 2 briefly reviews ABC and the SL. Sections 3 and 4 contain the details of our GP surrogate model and posterior estimation. Batch-sequential design strategies for sample-efficient estimation of the (approximate) posterior distribution are developed in Section 5 while Section 6 contains experiments. Finally, Section 7 contains discussion and concluding remarks. Proofs, implementation details and additional experiments can be found in the supplementary material (Järvenpää et al. 2020).

2 ABC and the synthetic likelihood methods

Our goal is to estimate parameters $\theta \in \Theta$ of a simulation model given observed data $\mathbf{x} \in \mathcal{X}$. We assume Θ is a compact subset of \mathbb{R}^d and that the prior information about feasible values of θ is coded into a (continuous) prior pdf $\pi(\theta)$. For simplicity we consider only continuous parameters but discrete parameters can be handled similarly. If evaluating the likelihood function $\pi(\mathbf{x} | \theta)$ is feasible, the posterior distribution can be computed using Bayes’ theorem $\pi(\theta | \mathbf{x}) \propto \pi(\theta)\pi(\mathbf{x} | \theta)$ up to a normalisation constant and hence be used as a target distribution in MCMC. However, when the likelihood is too costly to evaluate or unavailable, standard MCMC algorithms become infeasible.

Even when the likelihood is intractable, simulating “pseudo-data” from the model, i.e., drawing samples $\mathbf{x} \sim \pi(\cdot | \theta)$, is often feasible. In this case, ABC can be used for inference, see e.g. Marin et al. (2012); Turner and Van Zandt (2012); Lintusaari et al. (2017). Standard ABC techniques approximate the posterior as

$$\pi_{\text{ABC}}(\theta | \mathbf{x}) \propto \pi(\theta) \int_{\mathcal{X}} \mathbb{1}_{\Delta(\mathbf{x}, \mathbf{x}_s) \leq \varepsilon} \pi(\mathbf{x}_s | \theta) d\mathbf{x}_s, \quad (1)$$

where $\mathbb{1}$ denotes the indicator function, ε is a tolerance parameter and $\Delta : \mathcal{X}^2 \rightarrow \mathbb{R}_+$ is the discrepancy function used to compute the similarity between the simulated data \mathbf{x}_s and the observed data \mathbf{x} . The discrepancy is typically constructed from low-dimensional summary statistics $S : \mathcal{X} \rightarrow \mathbb{R}^p$, so that $\Delta(\mathbf{x}, \mathbf{x}_s) = \Delta'(S(\mathbf{x}), S(\mathbf{x}_s))$, where $\Delta' : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}_+$ is, for example, the weighted Euclidean distance. For each proposed parameter θ , an unbiased ABC posterior estimate can be obtained by replacing the integral in (1) with a Monte Carlo sum using some N simulated pseudo-data sets $\mathbf{x}^{(i)} \sim \pi(\cdot | \theta)$ for $i = 1, \dots, N$.

An alternative to ABC is the synthetic likelihood method (Wood, 2010; Price et al., 2018). In SL the summary statistics $S(\mathbf{x})$ are assumed to have a Gaussian distribution

for each parameter $\boldsymbol{\theta}$, that is

$$\pi(\mathbf{x}|\boldsymbol{\theta}) \approx \pi(S(\mathbf{x})|\boldsymbol{\theta}) \approx \pi_{\text{SL}}(S(\mathbf{x})|\boldsymbol{\theta}) \quad \mathcal{N}(S(\mathbf{x})|\boldsymbol{\mu}, \boldsymbol{\Sigma}). \quad (2)$$

The first approximation results from replacing the full data \mathbf{x} with a potentially non-sufficient summary statistics $S(\mathbf{x})$. The second approximation is due to the possible violations of the Gaussianity of $S(\mathbf{x})$. The expectation $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ in (2) are unknown and are estimated for each proposed parameter $\boldsymbol{\theta}$ using maximum likelihood (ML):

$$\hat{\boldsymbol{\mu}} = \frac{1}{N} \sum_{i=1}^N S(\mathbf{x}^{(i)}), \quad \hat{\boldsymbol{\Sigma}} = \frac{1}{N-1} \sum_{i=1}^N (S(\mathbf{x}^{(i)}) - \hat{\boldsymbol{\mu}})(S(\mathbf{x}^{(i)}) - \hat{\boldsymbol{\mu}})^\top, \quad (3)$$

where $\mathbf{x}^{(i)} \sim \pi(\cdot|\boldsymbol{\theta})$ for $i = 1, \dots, N$. As investigated by Price et al. (2018), the standard Metropolis algorithm can be combined with SL. The likelihood is then computed using (2) and the ML estimates in (3) or, alternatively, using an unbiased estimate of $\mathcal{N}(S(\mathbf{x})|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ shown in Section 2.1 of Price et al. (2018) which produces an exact pseudo-marginal MCMC if the Gaussianity assumption holds. See supplementary material C for discussion on the use of different SL estimators.

The advantage of SL over ABC is that specifying suitable ABC tuning parameters such as the tolerance and the discrepancy is avoided. While the Gaussianity of the summary statistics may not hold in practice, Price et al. (2018) have found that SL is often robust to deviations from normality. SL and its extensions (Thomas et al., 2018; An et al., 2019b,a; Frazier et al., 2019) produce pointwise noisy log-likelihood evaluations because in practice the number of repeated simulations N at each point is finite. Using (pseudo-marginal) MCMC or other sampling-based techniques for inference with these noisy targets thus requires a large number of simulations. Assuming noisy log-likelihood evaluations are available, e.g. obtained by using SL, the goal of the following sections is to develop an inference algorithm that can minimise the number of evaluations needed.

3 Gaussian process surrogate for the noisy log-likelihood

We denote the log-likelihood or its approximation, such as the log-SL obtained as the logarithm of (2), as $f(\boldsymbol{\theta}) = \log \pi(\mathbf{x}|\boldsymbol{\theta})$. We assume that we have access to noisy log-likelihood evaluations at $\boldsymbol{\theta}_i$ denoted by $y_i \in \mathbb{R}$ for building the surrogate model and that the “noise” i.e. the numerical or sampling error in evaluating the log-likelihood is independently Gaussian distributed. Treating the noisy log-likelihood evaluations y_i as “observations”, our measurement model is

$$y_i = f(\boldsymbol{\theta}_i) + \sigma_n(\boldsymbol{\theta}_i)\varepsilon_i, \quad \varepsilon_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad (4)$$

where $\sigma_n : \Theta \rightarrow (0, \infty)$ is a (continuous) function of $\boldsymbol{\theta}$ that determines the standard deviation of the observation noise and is assumed known. To justify our model in (4), we show empirically that log-SL is well approximated by a Gaussian distribution using six benchmark simulation models in the supplementary material D.1.

We place the following hierarchical GP prior for the log-likelihood function f :

$$f | \gamma \sim \mathcal{GP}(m_0(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad m_0(\boldsymbol{\theta}) = \sum_{i=1}^q \gamma_i h_i(\boldsymbol{\theta}), \quad \gamma \sim \mathcal{N}(\mathbf{b}, \mathbf{B}), \quad (5)$$

where $k : \Theta^2 \rightarrow \mathbb{R}$ is a covariance function and $h_i : \Theta \rightarrow \mathbb{R}$ are fixed basis functions (both assumed continuous). The nuisance parameters γ in (5) are marginalised, see e.g. O'Hagan and Kingman (1978); Rasmussen and Williams (2006), to obtain the following equivalent GP prior

$$f \sim \mathcal{GP}(\mathbf{h}(\boldsymbol{\theta})^\top \mathbf{b}, k(\boldsymbol{\theta}, \boldsymbol{\theta}') + \mathbf{h}(\boldsymbol{\theta})^\top \mathbf{B} \mathbf{h}(\boldsymbol{\theta}')), \quad (6)$$

where $\mathbf{h}(\boldsymbol{\theta}) \in \mathbb{R}^q$ is a column vector consisting of the basis functions h_i evaluated at $\boldsymbol{\theta}$. We use basis functions of the form $1, \theta_i, \theta_i^2$. A similar GP prior has been considered in Wilkinson (2014); Gutmann and Corander (2016); Drovandi et al. (2018), however, different from those articles, we take a fully Bayesian approach and marginalise γ as in Riihimäki and Vehtari (2014). Since little initial information is typically available on the magnitude and shape of the log-likelihood, we use relatively uninformative hyperpriors so that $\mathbf{b} = \mathbf{0}$ and $B_{ij} = 30^2 \mathbf{1}_{i=j}$. We assume that the log-likelihood function is smooth, and adopt the squared exponential covariance function $k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_f^2 \exp(-\sum_{i=1}^d (\theta_i - \theta'_i)^2 / (2l_i^2))$ although other choices, such as the Matérn covariance function, are also possible. We denote the $d+1$ covariance function hyperparameters as $\boldsymbol{\phi} = (\sigma_f^2, l_1, \dots, l_d)$. For now, we assume $\boldsymbol{\phi}$ is known and omit it from our notation for simplicity.

Given observations $D_t = \{(y_i, \boldsymbol{\theta}_i)\}_{i=1}^t$, which we call training data, our knowledge of the log-likelihood function is $f | D_t \sim \mathcal{GP}(m_t(\boldsymbol{\theta}), c_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$, where

$$m_t(\boldsymbol{\theta}) \quad k_t(\boldsymbol{\theta}) \mathbf{K}_t^{-1} \mathbf{y}_t + \mathbf{R}_t^\top(\boldsymbol{\theta}) \bar{\gamma}_t, \quad (7)$$

$$c_t(\boldsymbol{\theta}, \boldsymbol{\theta}') \quad k(\boldsymbol{\theta}, \boldsymbol{\theta}') - k_t(\boldsymbol{\theta}) \mathbf{K}_t^{-1} k_t^\top(\boldsymbol{\theta}') + \mathbf{R}_t^\top(\boldsymbol{\theta}) [\mathbf{B}^{-1} + \mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{H}_t^\top]^{-1} \mathbf{R}_t(\boldsymbol{\theta}'), \quad (8)$$

with $[\mathbf{K}_t]_{ij} \quad k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) + \mathbf{1}_{i=j} \sigma_n^2(\boldsymbol{\theta}_i)$ for $i, j \in \{1, \dots, t\}$, $k_t(\boldsymbol{\theta}) \quad (k(\boldsymbol{\theta}, \boldsymbol{\theta}_1), \dots, k(\boldsymbol{\theta}, \boldsymbol{\theta}_t))^\top$,

$$\bar{\gamma}_t \quad [\mathbf{B}^{-1} + \mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{H}_t^\top]^{-1} (\mathbf{H}_t \mathbf{K}_t^{-1} \mathbf{y}_t + \mathbf{B}^{-1} \mathbf{b}), \quad (9)$$

and $\mathbf{R}_t(\boldsymbol{\theta}) \quad \mathbf{H}(\boldsymbol{\theta}) - \mathbf{H}_t \mathbf{K}_t^{-1} k_t^\top(\boldsymbol{\theta})$. Above \mathbf{H}_t is the $q \times t$ matrix whose columns consist of basis function values evaluated at training points $\boldsymbol{\theta}_{1:t} = [\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_t]$ which is itself a $d \times t$ matrix, and $\mathbf{H}(\boldsymbol{\theta})$ is the corresponding $q \times 1$ vector at test point $\boldsymbol{\theta}$. From now on, we denote the GP variance function as $s_t^2(\boldsymbol{\theta}) \quad c_t(\boldsymbol{\theta}, \boldsymbol{\theta})$ and the probability law of f given D_t as $\Pi_{D_t}^f$, that is, $\Pi_{D_t}^f \quad \mathcal{GP}(m_t(\boldsymbol{\theta}), c_t(\boldsymbol{\theta}, \boldsymbol{\theta}'))$.

4 Estimators of the posterior from the GP surrogate

Using the GP surrogate model for the noisy log-likelihood, we here derive estimators for the posterior which can be e.g. plugged-in to a MCMC algorithm. Resulting sampling algorithms do not require further simulator runs (unlike e.g. SL-MCMC) producing potentially huge computational savings. Figure 1 demonstrates our approach. We want

to use our knowledge of the log-likelihood function represented by $\Pi_{D_t}^f$ to determine the optimal point estimate of the probability density function (pdf) of the posterior.¹ The uncertainty of the log-likelihood f can be propagated to the posterior distribution of the simulation model which consequently becomes a random quantity denoted as π_f :

$$\pi_f(\boldsymbol{\theta}) = \frac{\pi(\boldsymbol{\theta}) \exp(f(\boldsymbol{\theta}))}{\int_{\Theta} \pi(\boldsymbol{\theta}') \exp(f(\boldsymbol{\theta}')) d\boldsymbol{\theta}'}. \quad (10)$$

The expectation of the posterior pdf π_f at each parameter $\boldsymbol{\theta}$ can be formally written as

$$\mathbb{E}_{f|D_t}(\pi_f(\boldsymbol{\theta})) = \int \frac{\pi(\boldsymbol{\theta}) \exp(f(\boldsymbol{\theta}))}{\int_{\Theta} \pi(\boldsymbol{\theta}') \exp(f(\boldsymbol{\theta}')) d\boldsymbol{\theta}'} \Pi_{D_t}^f(df) \quad (11)$$

and the variance can be obtained similarly (assuming these quantities exist). In principle, one could sample posterior pdfs by first drawing $f^{(i)} \sim \Pi_{D_t}^f$ (a continuous function $\Theta \rightarrow \mathbb{R}$), then computing $\pi(\boldsymbol{\theta}) \exp(f^{(i)}(\boldsymbol{\theta}))$, and finally normalising. However, in practice this would require discretisation of the Θ -space and involves computational challenges. For this reason and similarly to Sinsbeck and Nowak (2017); Järvenpää et al. (2019), we instead take our quantity of interest to be the unnormalised posterior

$$\tilde{\pi}_f(\boldsymbol{\theta}) = \pi(\boldsymbol{\theta}) \exp(f(\boldsymbol{\theta})), \quad (12)$$

which follows log-Gaussian process that allows analytical computations.

Next we derive an optimal estimator for the unnormalised posterior $\tilde{\pi}$ in (12) using Bayesian decision-theory. We proceed here similarly to Sinsbeck and Nowak (2017) and consider the integrated quadratic loss function $l_2(\tilde{\pi}_1, \tilde{\pi}_2) = \int_{\Theta} (\tilde{\pi}_1(\boldsymbol{\theta}) - \tilde{\pi}_2(\boldsymbol{\theta}))^2 d\boldsymbol{\theta}$ between two (unnormalised) posterior densities $\tilde{\pi}_1$ and $\tilde{\pi}_2$. We assume $\tilde{\pi}_1$ and $\tilde{\pi}_2$ are square-integrable functions in Θ , i.e. $\tilde{\pi}_1, \tilde{\pi}_2 \in L^2(\Theta)$. The optimal Bayes estimator, denoted by $\hat{\tilde{\pi}} \in \mathbb{D}$, is the minimiser of the expected loss, where $\mathbb{D} = L^2(\Theta)$ denotes the set of candidate estimators. In detail,

$$\begin{aligned} \hat{\tilde{\pi}} &= \arg \min_{\tilde{d} \in \mathbb{D}} \mathbb{E}_{f|D_t} l_2(\tilde{\pi}_f, \tilde{d}) = \arg \min_{\tilde{d} \in \mathbb{D}} \mathbb{E}_{f|D_t} \int_{\Theta} (\tilde{\pi}_f(\boldsymbol{\theta}) - \tilde{d}(\boldsymbol{\theta}))^2 d\boldsymbol{\theta} \\ &= \arg \min_{\tilde{d} \in \mathbb{D}} \int_{\Theta} \mathbb{E}_{f|D_t} (\tilde{\pi}_f(\boldsymbol{\theta}) - \tilde{d}(\boldsymbol{\theta}))^2 d\boldsymbol{\theta}, \end{aligned} \quad (13)$$

where Tonelli theorem is used to change the order of expectation and integration and where $\tilde{d} \in \mathbb{D} = L^2(\Theta)$ is a candidate estimator of $\tilde{\pi}$. Equation 13 shows that the expected loss is minimised when the integrand on the second row is minimised independently for (almost) each $\boldsymbol{\theta} \in \Theta$. It follows from the basic results of Bayesian decision theory (see e.g. Robert 2007) that the minimum is obtained when $\tilde{d}(\boldsymbol{\theta}) = \mathbb{E}_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta}))$, i.e., the optimal estimator is the posterior expectation. The minimum value of (13), called

¹While in this article we are mainly concerned with point estimators of the posterior pdf, we can also quantify its (epistemic) uncertainty similarly to *probabilistic numerics* literature (see e.g. Hennig et al. 2015; Cockayne et al. 2019; Briol et al. 2019) as illustrated in Figure 1b. Such uncertainty estimates are also used to intelligently select the next simulation locations in Section 5.




Figure 1: (a) GP surrogate model for the log-SL of the Ricker model of Section 6.2 when only the first parameter, $\theta = \log(r)$, is varied. The black dots show the noisy log-SL evaluations and the black lines their approximate 95% confidence intervals, the grey area the 95% credible interval, and the red line the GP mean function. (b) Uncertainty of the SL. The grey area shows the 95% credible interval of the SL and the red line is the median estimate, obtained from (15). The dashed blue line shows the standard deviation of SL computed as the square root of (14).

Bayes risk, is the integrated variance $\int_{\Theta} \nabla_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) d\boldsymbol{\theta}$. The posterior expectation and variance can be computed from the log-Normal distribution as

$$\mathbb{E}_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_t(\cdot) + \frac{1}{2}s_t^2(\cdot)}, \quad \nabla_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi^2(\boldsymbol{\theta})e^{2m_t(\cdot) + s_t^2(\cdot)}(e^{s_t^2(\cdot)} - 1). \quad (14)$$

If we instead use L^1 loss $l_1(\tilde{\pi}_1, \tilde{\pi}_2) = \int_{\Theta} |\tilde{\pi}_1(\boldsymbol{\theta}) - \tilde{\pi}_2(\boldsymbol{\theta})| d\boldsymbol{\theta}$ where $\tilde{\pi}_1, \tilde{\pi}_2 \in L^1(\Theta)$, we can similarly show that the optimal point estimator is the marginal median. The median and the α -quantile q^α with $\alpha \in (0, 1)$ can be computed as

$$\text{med}_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_t(\cdot)}, \quad q_{f|D_t}^\alpha(\tilde{\pi}_f(\boldsymbol{\theta})) = \pi(\boldsymbol{\theta})e^{m_t(\cdot) + \Phi^{-1}(\alpha)s_t(\cdot)}, \quad (15)$$

where Φ^{-1} is the quantile function of the standard normal distribution. As we show in the supplementary material A, the Bayes risk corresponding to the L^1 loss is

$$\min_{\tilde{d} \in \mathcal{D}} \mathbb{E}_{f|D_t} l_1(\tilde{\pi}_f, \tilde{d}) = \int_{\Theta} \pi(\boldsymbol{\theta}) \exp(m_t(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta})/2) (2\Phi(s_t(\boldsymbol{\theta})) - 1) d\boldsymbol{\theta}. \quad (16)$$

When MCMC is used with the point estimator of the unnormalised posterior in either (14) or (15), we are in fact targeting the following mean and median based estimators of the (normalised) posterior

$$\pi_t^{\text{mean}}(\boldsymbol{\theta}) = \frac{\pi(\boldsymbol{\theta})e^{m_t(\cdot) + \frac{1}{2}s_t^2(\cdot)}}{\int_{\Theta} \pi(\boldsymbol{\theta}')e^{m_t(\cdot) + \frac{1}{2}s_t^2(\cdot)} d\boldsymbol{\theta}'}, \quad \pi_t^{\text{med}}(\boldsymbol{\theta}) = \frac{\pi(\boldsymbol{\theta})e^{m_t(\cdot)}}{\int_{\Theta} \pi(\boldsymbol{\theta}')e^{m_t(\cdot)} d\boldsymbol{\theta}'}. \quad (17)$$

These are obtained by simply normalising the Bayes optimal estimators of the unnormalised posterior (and, as a consequence, a guarantee of optimality for normalised

posterior does not follow). Similar estimators were also considered by Stuart and Teckentrup (2018). Both are clearly valid density functions, and tractable, unlike (11). The latter, i.e., the marginal median based estimate, is equal to (10) if we replace the unknown log-likelihood function $f(\boldsymbol{\theta})$ with a GP mean function $m_t(\boldsymbol{\theta})$ and neglect GP uncertainty. On the other hand, the former, i.e., the marginal mean estimate, takes into account the GP uncertainty through the variance function $s_t^2(\boldsymbol{\theta})$. These two point estimates become the same if the GP variance is negligible.

5 Parallel designs of simulations

In the previous section we showed how to quantify the uncertainty of the GP surrogate-based unnormalised posterior density and we derived computable and (in a certain sense) optimal point estimates of it. Next we develop Bayesian experimental design strategies to select further locations to evaluate the log-likelihood, so that the uncertainty in the unnormalised posterior decreases as fast as possible. We focus on batch strategies and denote the batch size as $b \in \{1, 2, \dots\}$. Before moving on, we introduce some terminology. The next batch of b evaluation locations is obtained as the solution to an optimisation problem. We call the objective function of this optimisation problem a *design criterion* and the resulting batch of evaluation locations as *design points* or just *design*. The complete procedure of selecting the design points is called a *batch-sequential* (or, when $b = 1$, just *sequential*) *strategy*.²

In this paper we focus on synchronous parallelisation where a batch of b design points is constructed at each iteration and the corresponding b simulations are simultaneously submitted to the workers. However, the “greedy” design strategies developed in Sections 5.3 and 5.5 can also be used for asynchronous parallelisation, where a new location is immediately chosen and submitted for processing, whenever any of the running simulations completes, instead of waiting all the other $b - 1$ simulations to finish.

5.1 Analytical expressions for the design criteria

We first derive some general results needed for efficient evaluation of the design criteria. These can be useful also for developing batch designs for other related GP-based problems such as BQ and BO. Given $D_t = \{(y_i, \boldsymbol{\theta}_i)\}_{i=1}^t$ it is useful to know how additional b candidate evaluations at points $\boldsymbol{\theta}^* = [\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_b^*]$ would affect our knowledge about the log-likelihood f and the unnormalised posterior $\tilde{\pi}$. The following Lemma is central to our analysis. It shows how the GP mean and variance functions are affected by supplementing the training data D_t with a new batch of evaluations $D^* = \{(y_i^*, \boldsymbol{\theta}_i^*)\}_{i=1}^b$ when the unknown \mathbf{y}^* is assumed to be distributed according to the posterior predictive distribution of the GP given D_t . The Lemma is a generalisation of a similar result by Järvenpää et al. (2019); Lyu et al. (2018).

Lemma 5.1. *Consider the mean and variance functions of the GP model in Section 3 for a fixed $\boldsymbol{\theta}$, given the training data $D_t \cup D^*$ and when treated as functions*

²The design criterion is often called an acquisition function and the resulting strategy sometimes an acquisition rule in the BO literature.

of \mathbf{y}^* . Assume \mathbf{y}^* follows the posterior predictive distribution, that is $\mathbf{y}^* | \boldsymbol{\theta}^*, D_t \sim \mathcal{N}(m_t(\boldsymbol{\theta}^*), c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \text{diag}(\sigma_n^2(\boldsymbol{\theta}_1^*), \dots, \sigma_n^2(\boldsymbol{\theta}_b^*)))$. Then,

$$m_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*) | \boldsymbol{\theta}^*, D_t \sim \mathcal{N}(m_t(\boldsymbol{\theta}), \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)), \quad (18)$$

$$s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) | \boldsymbol{\theta}^*, D_t \sim \delta(s_t^2(\boldsymbol{\theta}) - \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) - s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)), \quad (19)$$

where $\delta(\cdot)$ is the Dirac measure and

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = c_t(\boldsymbol{\theta}, \boldsymbol{\theta}^*) [c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) + \text{diag}(\sigma_n^2(\boldsymbol{\theta}_1^*), \dots, \sigma_n^2(\boldsymbol{\theta}_b^*))]^{-1} c_t(\boldsymbol{\theta}^*, \boldsymbol{\theta}). \quad (20)$$

In the Lemma, $m_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ is the GP mean function at iteration $t + b$ whose dependence on $\boldsymbol{\theta}^*$ is shown explicitly. Importantly, the above Lemma shows how the GP variance decreases from $s_t^2(\boldsymbol{\theta})$ to $s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ when the extra b evaluations at $\boldsymbol{\theta}^*$ are included, and the reduction $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ is deterministic. We see, for example, that if $c_t(\boldsymbol{\theta}_i^*, \boldsymbol{\theta}_j^*) = 0$ for all $i, j = 1, \dots, b, i \neq j$ which might hold approximately, e.g., if the evaluation points $\boldsymbol{\theta}_i^*$ are located far from each other, then

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \sum_{i=1}^b \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_i^*) = \sum_{i=1}^b \frac{c_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}_i^*)}{s_t^2(\boldsymbol{\theta}_i^*) + \sigma_n^2(\boldsymbol{\theta}_i^*)}. \quad (21)$$

This shows that the reduction of GP variance at $\boldsymbol{\theta}$, $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$, factorises over the new evaluation points $\boldsymbol{\theta}_i^*$ in $\boldsymbol{\theta}^*$. Intuitively, if the test point $\boldsymbol{\theta}$ is strongly correlated with some evaluation point $\boldsymbol{\theta}_i^*$, including the evaluation at $\boldsymbol{\theta}_i^*$ will result in a large reduction of variance at the test point. Furthermore, the larger the noise variance $\sigma_n^2(\boldsymbol{\theta}_i^*)$ at the evaluation point $\boldsymbol{\theta}_i^*$ is, the less the GP variance will decrease.

It clearly holds that $0 \leq \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) \leq s_t^2(\boldsymbol{\theta})$. Items (i–ii) of the following Lemma summarise some additional properties of the variance reduction function in (20) and (iii–iv) show two further useful identities needed later. Item (i) shows the (rather obvious) result that the order of evaluation points in $\boldsymbol{\theta}^*$ does not change $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ and in the following we often identify the $d \times b$ matrix $\boldsymbol{\theta}^*$ with a multiset whose elements are the columns of $\boldsymbol{\theta}^*$ although this leads to some abuse of notation.

Lemma 5.2. *Let $\boldsymbol{\theta}^* \in \mathbb{R}^{d \times b}$ and let $\boldsymbol{\theta} \in \mathbb{R}^d$ be any test point. The function $\boldsymbol{\theta}^* \mapsto \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ in (20) for any fixed $\boldsymbol{\theta}$ has the following properties.*

- (i) *The function value is invariant to the permutation of the evaluation locations i.e. the columns of $\boldsymbol{\theta}^*$.*
- (ii) *Let $\boldsymbol{\theta}_A^* \subseteq \boldsymbol{\theta}_B^*$. Then $\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_A^*) \leq \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_B^*)$, i.e., including new evaluations never increases variance.*
- (iii) *Let $\boldsymbol{\theta}^* = [\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*]$ so that $b = 2$ and denote the predictive variance at $\boldsymbol{\theta}_j^*$ by $\bar{s}_t^2(\boldsymbol{\theta}_j^*) = s_t^2(\boldsymbol{\theta}_j^*) + \sigma_n^2(\boldsymbol{\theta}_j^*)$ for $j \in \{1, 2\}$. Then*

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_1^*) + \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_2^*) + r_t(\boldsymbol{\theta}; \boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*), \quad (22)$$

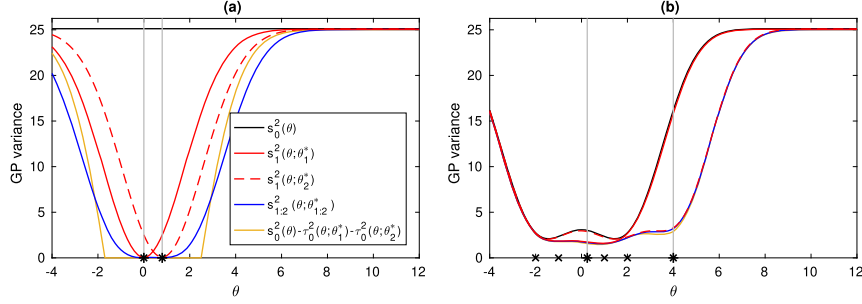


Figure 2: The effect of two new evaluations (black stars) on the GP variance. Black line is the original variance, red lines (solid and dashed) show variance if only one of the evaluations is included. Blue line shows the variance after both evaluations, and yellow the variance if the interaction between the locations is neglected. (a) Noiseless observations at evaluation locations close to each other are obtained. (b) similar to (a), but showing four earlier evaluations (at black crosses) from which noisy observations were available, such that the GP variance is not exactly zero at these locations.

$$r_t(\boldsymbol{\theta}; \boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*) = \frac{c_t^2(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)c_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}_1^*)\bar{s}_t^2(\boldsymbol{\theta}_2^*) + c_t^2(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)c_t^2(\boldsymbol{\theta}, \boldsymbol{\theta}_2^*)\bar{s}_t^2(\boldsymbol{\theta}_1^*)}{\bar{s}_t^4(\boldsymbol{\theta}_1^*)\bar{s}_t^4(\boldsymbol{\theta}_2^*) - \bar{s}_t^2(\boldsymbol{\theta}_1^*)\bar{s}_t^2(\boldsymbol{\theta}_2^*)c_t^2(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)} - 2 \frac{c_t(\boldsymbol{\theta}, \boldsymbol{\theta}_1^*)c_t(\boldsymbol{\theta}, \boldsymbol{\theta}_2^*)c_t(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)}{\bar{s}_t^2(\boldsymbol{\theta}_1^*)\bar{s}_t^2(\boldsymbol{\theta}_2^*) - c_t^2(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)}. \quad (23)$$

(iv) Let $\boldsymbol{\theta}^* = [\boldsymbol{\theta}_A^*, \boldsymbol{\theta}_b^*]$ where $\boldsymbol{\theta}_A^* \in \mathbb{R}^{d \times (b-1)}$ and $\boldsymbol{\theta}_b^* \in \mathbb{R}^d$, and denote $\bar{\mathbf{S}}_A = c_t(\boldsymbol{\theta}_A^*) + \text{diag}(\sigma_n^2(\boldsymbol{\theta}_1^*), \dots, \sigma_n^2(\boldsymbol{\theta}_{b-1}^*))$. Then

$$\tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) = \tau_t^2(\boldsymbol{\theta}; \boldsymbol{\theta}_A^*) + \frac{(c_t(\boldsymbol{\theta}, \boldsymbol{\theta}_b^*) - c_t(\boldsymbol{\theta}, \boldsymbol{\theta}_A^*)\bar{\mathbf{S}}_A^{-1}c_t(\boldsymbol{\theta}_A^*, \boldsymbol{\theta}_b^*))^2}{\bar{s}_t^2(\boldsymbol{\theta}_b^*) + \sigma_n^2(\boldsymbol{\theta}_b^*) - c_t(\boldsymbol{\theta}_b^*, \boldsymbol{\theta}_A^*)\bar{\mathbf{S}}_A^{-1}c_t(\boldsymbol{\theta}_A^*, \boldsymbol{\theta}_b^*)}. \quad (24)$$

Figure 2 demonstrates how two new evaluations at $\boldsymbol{\theta}_1^*$ and $\boldsymbol{\theta}_2^*$ reduce the GP variance in two different one-dimensional examples. Specifically, Figure 2a illustrates the fact of Lemma 5.2 (iii) that the interaction between the evaluation points, represented by the term $r_t(\boldsymbol{\theta}; \boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)$, affects the reduction of the GP variance and its effect can be either positive or negative (the variance after two new evaluations, blue line, is either below or above the yellow line, which represents the reduced variance if the interaction is neglected). In Figure 2b the new evaluation locations are far apart and the factorisation of the variance reduction in (21) holds approximately.

5.2 Batch-sequential designs

Given $D_t = \{(y_i, \boldsymbol{\theta}_i)\}_{i=1}^t$, our goal is to select the next batch of b evaluations $\boldsymbol{\theta}^*$ in an optimal fashion. We take a Bayesian decision theoretic approach, where $\boldsymbol{\theta}^*$ is selected to minimise the expected (or median) loss, where the loss measures uncertainty remaining

in the unnormalised posterior $\tilde{\pi}^f$ when the hypothetical observations \mathbf{y}^* at locations $\boldsymbol{\theta}^*$ are taken into account. In the following we develop two such techniques based on two different measures of uncertainty: variance and interquartile range (IQR). Design strategies which acknowledge the impact of the next batch, but neglect the whole remaining computational budget, are often called “myopic”. It is possible to formulate a non-myopic design as a dynamic programming problem, but this is computationally demanding, see e.g. Bect et al. (2012); González et al. (2016). Consequently, we focus on myopic designs which already produce highly sample-efficient and practical algorithms.

Expected integrated variance (EIV)

As our first measure of the uncertainty of the unnormalised posterior $\tilde{\pi}^f$ for the selection of the next batch design $\boldsymbol{\theta}^*$, we select the Bayes risk under the L^2 loss. In this case, the Bayes risk is the integrated variance function

$$\mathcal{L}^v(\Pi_{D_t}^f) = \int_{\Theta} \mathbb{V}_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) d\boldsymbol{\theta} = \int_{\Theta} \pi^2(\boldsymbol{\theta}) e^{2m_t(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta})} \left(e^{s_t^2(\boldsymbol{\theta})} - 1 \right) d\boldsymbol{\theta}, \quad (25)$$

whose integrand was obtained from (14). This is similar to Sinsbeck and Nowak (2017); Järvenpää et al. (2019) who, however, considered other GP surrogate models and only sequential designs. We compute the expectation over the hypothetical noisy log-likelihoods \mathbf{y}^* for any candidate design $\boldsymbol{\theta}^*$, leading to the expected integrated variance design criterion (EIV). The resulting optimal strategy is a special case of stepwise uncertainty reduction technique, see e.g. Bect et al. (2012). This criterion is evaluated efficiently without numerical simulations from the GP model, using the following result.

Proposition 5.3. *With the assumptions of Lemma 5.1, the expected integrated variance design criterion L_t^v at any candidate design $\boldsymbol{\theta}^* \in \Theta^b$ is*

$$L_t^v(\boldsymbol{\theta}^*) = \mathbb{E}_{\mathbf{y}^* | \boldsymbol{\theta}^*} \mathcal{L}^v(\Pi_{D_t \cup D}^f) = \int_{\Theta} \pi^2(\boldsymbol{\theta}) e^{2m_t(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta})} \left(e^{s_t^2(\boldsymbol{\theta})} - e^{\tau_t^2(\boldsymbol{\theta})} \right) d\boldsymbol{\theta}. \quad (26)$$

Integrated median interquartile range (IMIQR)

A sequential design strategy based on EIV worked well in the ABC scenario of Järvenpää et al. (2019), who modelled the discrepancy in (1) with a GP. In this article we instead model the log-likelihood with a GP as illustrated in Figure 1 and the goal is to minimise the uncertainty of the unnormalised posterior $\tilde{\pi}^f$, which has a log-Normal distribution for a fixed $\boldsymbol{\theta}$. However, the expectation and variance can be suboptimal estimates of the central tendency and uncertainty of a heavy-tailed distribution such as log-Normal. For example, Figure 1b shows that the standard deviation (dashed blue line) grows very rapidly at the boundaries although at the same time the credible interval clearly indicates that the probability of the log-likelihood, and consequently the likelihood, of having a non-negligible value there is vanishingly small. The mean is similarly affected in a non-intuitive way by the heavy tails: It is fairly easy to see that

$$\mathbb{P}[\exp(f(\boldsymbol{\theta})) \geq \mathbb{E}(\exp(f(\boldsymbol{\theta})))] = \mathbb{P}(f(\boldsymbol{\theta}) \geq m_t(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta})/2) = \Phi(-s_t(\boldsymbol{\theta})/2). \quad (27)$$

This means that with a sufficiently large variance of the log-likelihood $s_t^2(\boldsymbol{\theta})$, the probability that the likelihood $\exp(f(\boldsymbol{\theta}))$ is greater than its own mean becomes negligible.

The above analysis suggests (and empirical results in Section 6 further confirm) that mean-based point estimates and variance-based design strategies, such as the EIV and those proposed by Gunter et al. (2014); Kandasamy et al. (2015); Sinsbeck and Nowak (2017); Järvenpää et al. (2019); Acerbi (2018), are unsuitable when log-likelihood evaluations are noisy. A reasonable alternative for the L^2 -loss used to derive the EIV is to measure the uncertainty in the posterior using the L^1 -loss, which is less affected by extreme values. As shown in Section 3, the L^1 -loss leads to the marginal median estimate for the posterior, π_t^{med} . While the L^1 -loss produces a robust median estimator that we adopt, (16) shows that the Bayes risk with L^1 loss scales as $\exp(s_t^2(\boldsymbol{\theta})/2)$ since $\Phi(s_t(\boldsymbol{\theta})) \approx 1$ for large $s_t(\boldsymbol{\theta})$, such that also this measure for overall uncertainty of $\tilde{\pi}^f$ is affected by the heavy tails of the log-Normal distribution.

We propose a new, robust design criterion for selecting the next design. In place of the variance in EIV, we use a robust measure of uncertainty, the interquartile range $\text{IQR}(\boldsymbol{\theta}) = q^{3/4}(\boldsymbol{\theta}) - q^{1/4}(\boldsymbol{\theta})$. The integrated IQR loss measuring the uncertainty of the posterior pdf is defined as

$$\mathcal{L}^{\text{IQR}}(\Pi_{D_t}^f) = \int_{\Theta} \text{IQR}_{f|D_t}(\tilde{\pi}_f(\boldsymbol{\theta})) \, d\boldsymbol{\theta} = 2 \int_{\Theta} \pi(\boldsymbol{\theta}) e^{m_t(\boldsymbol{\theta})} \sinh(us_t(\boldsymbol{\theta})) \, d\boldsymbol{\theta}, \quad (28)$$

where $u = \Phi^{-1}(p_u)$ and $\sinh(z) = (\exp(z) - \exp(-z))/2$ for $z \in \mathbb{R}$ is the hyperbolic sine, which emerges after using (15). While we use $p_u = 0.75$, other quantiles $p_u \in (0.5, 1)$ are also possible. A theoretical downside of the IQR loss is that it does not formally coincide with the Bayes risk for the L^1 or L^2 loss, which correspond to the optimal point estimators of the unnormalised posterior (see Section 4).

We also use the median in place of the mean to measure the effect of the next design $\boldsymbol{\theta}^*$ to the loss function. That is, we use median loss decision theory (see Yu and Clarke 2011), and define the median integrated IQR loss function as

$$L_t^{\text{IQR}}(\boldsymbol{\theta}^*) = \text{med}_{\mathbf{y} | \cdot, D_t} \mathcal{L}^{\text{IQR}}(\Pi_{D_t \cup D}^f). \quad (29)$$

The median integrated IQR loss in (29) is intractable but it can be approximated by the integrated median IQR loss (IMIQR). This approximation³ follows by replacing the predictive distribution of \mathbf{y}^* with a point mass, i.e., $\pi(\mathbf{y}^* | \boldsymbol{\theta}^*, D_t) \approx \delta(m_t(\boldsymbol{\theta}^*) - \mathbf{y}^*)$. This approximation resembles the so-called kriging believer heuristic in Ginsbourger et al. (2010). The next result gives a useful formula to calculate IMIQR.

³Instead of approximation, which may be inaccurate when the GP variance function is large, the integrated median criterion in (30) can be seen as an alternative decision-theoretic formulation with infinitely many dependent variables of interest (one for each $\boldsymbol{\theta} \in \Theta$) and where the median outcomes after considering the effect of the design $\boldsymbol{\theta}$ are all computed separately for each $\boldsymbol{\theta} \in \Theta$, and the corresponding losses are combined through averaging. The median integrated loss in (29) instead has a single combined loss function.

Proposition 5.4. *With the assumptions of Lemma 5.1, the integrated median IQR loss, denoted as \tilde{L}_t^{IQR} , at any candidate design $\boldsymbol{\theta}^* \in \Theta^b$ is*

$$\tilde{L}_t^{\text{IQR}}(\boldsymbol{\theta}^*) = \int_{\Theta} \text{med}_{\mathbf{y} | \cdot, D_t} \text{IQR}_{f|D_t \cup D}(\tilde{\pi}_f(\boldsymbol{\theta})) d\boldsymbol{\theta} = 2 \int_{\Theta} \pi(\boldsymbol{\theta}) e^{m_t(\cdot)} \sinh(us_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*)) d\boldsymbol{\theta}. \quad (30)$$

The integrand of (30) is recognised as a product of the marginal median estimate of the posterior in (15) and the function $\sinh(us_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*))$. Hence, to minimise IMIQR, the simulation locations $\boldsymbol{\theta}^*$ need to be chosen as a compromise between regions where the current posterior estimate is non-negligible and where the GP variance $s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ decreases efficiently when the simulations are run at $\boldsymbol{\theta}^*$. Similar interpretation holds also for the EIV function in (26). However, EIV assigns significantly more weight to areas with high GP variance than IMIQR.

5.3 Joint and greedy optimisation for batch-sequential designs

We can now evaluate EIV and IMIQR design criteria for any candidate design $\boldsymbol{\theta}^*$ and choose $\boldsymbol{\theta}^*$ as the minimiser, i.e.,

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta} \in \Theta^b} L_t(\boldsymbol{\theta}), \quad (31)$$

where L_t is either the EIV in (26) or IMIQR in (30). The objective function is typically smooth but multimodal so global optimisation is needed. We call (31) as “joint” optimisation which does not scale to high dimensional parameter spaces or to large batch sizes. Even if computing the design criterion is cheap as compared to the run times of typical simulation models, solving the db -dimensional global optimisation problem is often impractical as discussed in Wilson et al. (2018). Hence, we consider greedy optimisation as also used in batch BO (Ginsbourger et al., 2010; Snoek et al., 2012; Wilson et al., 2018). The greedy optimisation procedure for both EIV and IMIQR works as follows: the first point $\boldsymbol{\theta}_1^*$ is chosen as in the sequential case i.e. by solving (31) with $b = 1$. The rest of the points $\boldsymbol{\theta}_{2:b}^*$ are obtained by iteratively solving

$$\boldsymbol{\theta}_r^* = \arg \min_{\boldsymbol{\theta} \in \Theta} L_t([\boldsymbol{\theta}_{1:r-1}^*, \boldsymbol{\theta}]), \quad r = 2, 3, \dots, b. \quad (32)$$

This greedy approach simplifies the difficult db -dimensional optimisation into b separate d -dimensional problems, and makes it scalable as a function of b .

In general, the design found by the greedy optimisation does not equal the minimiser of the joint criterion. It follows from Lemma 5.2 (i) that both EIV and IMIQR are invariant to the order of evaluation locations in $\boldsymbol{\theta}^*$ but this does not hold for the greedy procedure. Bounds for the performance of greedy maximisation of a set function have been studied in literature, see e.g. Nemhauser et al. (1978); Krause et al. (2008); Bach (2013). For example, if the design criterion (when defined equivalently using a utility so that (32) becomes a maximisation problem) is submodular and non-decreasing in batch size b , then the worst-case outcome of greedy optimisation is at least $1 - 1/e \approx 0.63$

of the corresponding optimal joint value. A utility function corresponding to IMIQR defined below is not submodular but an approximation of it is weakly submodular (see e.g. Krause et al. 2008; Krause and Cevher 2010). We use this fact to derive a weaker but still useful bound.

We here consider an approximation $\tilde{L}_t^{\text{IQR},a}(\boldsymbol{\theta}^*)$ of $\tilde{L}_t^{\text{IQR}}(\boldsymbol{\theta}^*)$ so that

$$\tilde{L}_t^{\text{IQR}}(\boldsymbol{\theta}^*) \approx \tilde{L}_t^{\text{IQR},a}(\boldsymbol{\theta}^*) - 2u^2 \int_{\Theta} \pi(\boldsymbol{\theta}) e^{m_t(\cdot)} s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*) d\boldsymbol{\theta}, \quad (33)$$

which follows from the observation that $\sinh(us_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*)) \approx u^2 s_{t+b}^2(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$ and where we had $u = \Phi^{-1}(p_u)$. The approximation in (33) is reasonable when $s_{t+b}(\boldsymbol{\theta}; \boldsymbol{\theta}^*) \in [0, 3/u]$ in the region where $\pi(\boldsymbol{\theta}) e^{m_t(\cdot)}$ is non-negligible. For simplicity, we consider a discretised setting where the optimisation is done over a finite set $\tilde{\Theta} \subset \Theta$ and define (approximate) IMIQR utility function as

$$\tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta}^*) = \tilde{L}_t^{\text{IQR},a}(\emptyset) - \tilde{L}_t^{\text{IQR},a}(\boldsymbol{\theta}^*) \quad (34)$$

for $\boldsymbol{\theta}^* \in 2\tilde{\Theta}$. Clearly, maximising $\tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta}^*)$ is equivalent to minimising $\tilde{L}_t^{\text{IQR},a}(\boldsymbol{\theta}^*)$. The following theorem gives a bound for the greedy optimisation of the (approximate) IMIQR utility function. The proof can be found in supplementary material B.2.

Theorem 5.5. *Consider the set function $\tilde{U}_t^{\text{IQR},a} : 2\tilde{\Theta} \rightarrow \mathbb{R}_+$ in (34). Let $\boldsymbol{\theta}_O$ be a (joint) optimal solution for maximising $\tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta})$ over $\boldsymbol{\theta} \subset \tilde{\Theta}, |\boldsymbol{\theta}| \leq b$. The greedy algorithm for this maximisation problem outputs a set $\boldsymbol{\theta}_G \subset \tilde{\Theta}$ satisfying*

$$\tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta}_G) \geq (1 - 1/e) \tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta}_O) - b^2 \varepsilon_t, \quad \text{where} \quad (35)$$

$$\varepsilon_t = \max\{0, 2u^2 \varepsilon'_t\}, \quad \varepsilon'_t = \max_{\substack{A \subset \tilde{\Theta}, |A|=i \leq 2b \\ j, k \in \tilde{\Theta}}} \int_{\Theta} \pi(\boldsymbol{\theta}) e^{m_t(\cdot)} r_{1:t+i}(\boldsymbol{\theta}; \boldsymbol{\theta}_j, \boldsymbol{\theta}_k) d\boldsymbol{\theta}. \quad (36)$$

Computing ε_t explicitly is difficult but we expect that often $\varepsilon_t \ll \tilde{U}_t^{\text{IQR},a}(\boldsymbol{\theta}_O)$ since $r_{1:t+i}(\boldsymbol{\theta}; \boldsymbol{\theta}_j, \boldsymbol{\theta}_k)$ given by (23) tends to be small. However, ε_t may not always be small, the term $b^2 \varepsilon_t$ scales quadratically for batch size b , and the bound holds only approximately for IMIQR. This bound still suggests that, at least in some iterations of the algorithm, greedy IMIQR produces near-optimal batch locations. On the other hand, an approximation similar to (33) for EIV would be reasonable in a very limited number of situations, and experiments in supplementary material D.3 suggest that greedy EIV scales worse as a function of b compared to the corresponding greedy IMIQR strategy. Finally, we note that even when the bound is weak, new design points cannot increase the value of EIV or IMIQR loss function as shown in supplementary material B.1. Hence, the batch strategies cannot be worse than the corresponding sequential designs and, in practice, they are highly useful as is seen empirically in Section 6.

5.4 Implementation details

Using the GP surrogate model and the analysis from the previous sections, we show the resulting inference method as Algorithm 1. Some key implementation details are

discussed below and further details (e.g. on handling GP hyperparameters, MCMC methods used, and optimisation of the design criteria) are given in supplementary material C. The algorithm is shown for the SL case using the IMIQR strategy, but it works similarly for EIV, heuristic designs developed in the next section and other log-likelihood estimators besides SL. The potentially expensive simulations on the lines 2–5 and 18–21 can be done in parallel. In the SL case, the simulations can be parallelised in terms of both the number of repeated simulations N and batch size b .

Algorithm 1 GP-based SL inference using IMIQR with synchronous batch design.

Input: Prior density $\pi(\boldsymbol{\theta})$, simulation model $\pi(\cdot | \boldsymbol{\theta})$, GP prior Π^f , number of repeated samples N , summary function S , batch size b , initial batch size b_0 , max. iterations i_{\max} , number of IS samples s_{IS} , number of MCMC samples s_{MC}

- 1: Sample $\boldsymbol{\theta}_{1:b_0} \stackrel{\text{i.i.d.}}{\sim} \pi(\cdot)$ (or use some other space-filling initial design)
- 2: **for** $r = 1 : b_0$ **do**
- 3: Simulate $\mathbf{x}_r^{(1:N)} \stackrel{\text{i.i.d.}}{\sim} \pi(\cdot | \boldsymbol{\theta}_r)$ and compute $S_r^{(1:N)} = S(\mathbf{x}_r^{(1:N)})$
- 4: Compute y_r from $\{S_r^{(j)}\}_{j=1}^N$
- 5: **end for**
- 6: Set initial training data $D_{b_0} \leftarrow \{(y_r, \boldsymbol{\theta}_r)\}_{r=1}^{b_0}$
- 7: **for** $i = 1 : i_{\max}$ **do**
- 8: Use MAP estimation to obtain GP hyperparameters $\boldsymbol{\phi}$ using $D_{b_0+(i-1)b}$
- 9: Sample $\boldsymbol{\theta}^{(j)} \sim \pi_q$ using MCMC and compute $\omega^{(j)}$ in Eq. 37 for $j = 1, \dots, s_{\text{IS}}$
- 10: **if** `joint_optim` **then**
- 11: Obtain $\boldsymbol{\theta}_{1:b}^{(i)*}$ by solving Eq. 31 using Eq. 30 and 37
- 12: **else if** `greedy_optim` **then**
- 13: Obtain $\boldsymbol{\theta}_1^{(i)*}$ by minimising Eq. 31 using Eq. 30 and 37
- 14: **for** $r = 2 : b$ **do**
- 15: Obtain $\boldsymbol{\theta}_r^{(i)*}$ by solving Eq. 32 using Eq. 30 and 37
- 16: **end for**
- 17: **end if**
- 18: **for** $r = 1 : b$ **do**
- 19: Simulate $\mathbf{x}_r^{(i,1:N)} \stackrel{\text{i.i.d.}}{\sim} \pi(\cdot | \boldsymbol{\theta}_r^{(i)*})$, compute $S_r^{(i,1:N)} = S(\mathbf{x}_r^{(i,1:N)})$
- 20: Compute $y_r^{(i)*}$ using $\{S_r^{(i,j)}\}_{j=1}^N$
- 21: **end for**
- 22: Update training data $D_{b_0+ib} \leftarrow D_{b_0+(i-1)b} \cup \{(y_r^{(i)*}, \boldsymbol{\theta}_r^{(i)*})\}_{r=1}^b$
- 23: **end for**
- 24: Use MAP estimation to obtain GP hyperparameters $\boldsymbol{\phi}$ using $D_{b_0+i_{\max}b}$
- 25: Sample $\boldsymbol{\vartheta}^{(1:s_{\text{MC}})}$ from the marginal median estimate in Eq. 17 using MCMC
- 26: **return** Samples $\boldsymbol{\vartheta}^{(1:s_{\text{MC}})}$ from the approximate SL posterior

Evaluation of EIV and IMIQR requires numerical integration over $\Theta \subset \mathbb{R}^d$. Similar computational challenges emerge also in the state-of-the-art BO methods such as Hennig and Schuler (2012); Hernández-Lobato et al. (2014); Wu and Frazier (2016) and in Chevalier et al. (2014). If $d \leq 2$ we discretise the parameter space Θ and approximate the integral in the resulting grid. In higher dimensions, we use self-normalised importance

sampling (IS) as in Chevalier et al. (2014); Järvenpää et al. (2019). Specifically, we draw samples from the importance distribution $\boldsymbol{\theta}^{(j)} \sim \pi_q(\boldsymbol{\theta})$ and use these as integration points to approximate

$$\int_{\Theta} I_t(\boldsymbol{\theta}; \boldsymbol{\theta}^*) d\boldsymbol{\theta} \approx \sum_{j=1}^{\text{sis}} \omega^{(j)} I_t(\boldsymbol{\theta}^{(j)}; \boldsymbol{\theta}^*), \quad \omega^{(j)} = \frac{1/\pi_q(\boldsymbol{\theta}^{(j)})}{\sum_{k=1}^s 1/\pi_q(\boldsymbol{\theta}^{(k)})}, \quad (37)$$

where the integrand of either (26) or 30 is denoted by $I_t(\boldsymbol{\theta}; \boldsymbol{\theta}^*)$. As the proposal π_q we use the current loss (see (14) and the integrand of (16)), which is a function $\Theta \rightarrow \mathbb{R}_+$, and can be interpreted as an unnormalised pdf. This is a natural choice because the current loss has a similar shape as the expected/median loss as a function of $\boldsymbol{\theta}$. We use the same proposal in the greedy optimisation in (32) although it would be also possible to adapt the proposal π_q according to the pending points $\boldsymbol{\theta}_{1:r-1}^*$ when optimising with respect to the r th point $\boldsymbol{\theta}_r^*$.

We have assumed that the noise function σ_n^2 in (4) is known. In practice, this is a valid assumption only in the noiseless case where $\sigma_n^2(\boldsymbol{\theta}) = 0$. As our focus is on the noisy setting, we need to estimate σ_n^2 . Sometimes σ_n^2 can be assumed to be an unknown constant to be determined together with the GP hyperparameters $\boldsymbol{\phi}$ using MAP estimation. However, we observed that σ_n^2 often depends on the magnitude of the log-likelihood (see Figure 1), making the assumption of homoscedastic noise questionable. Similarly to Wilkinson (2014), we estimate σ_n^2 using the bootstrap. Specifically, with each new training data point $\boldsymbol{\theta}_i$, we resample with replacement N summary vectors from the original population $\{S_i^{(j)}\}_{j=1}^N$ for 2000 times. We then compute the empirical variance of the resulting log-SL values and use it as a plug-in estimator for $\sigma_n^2(\boldsymbol{\theta}_i)$.

For EIV, IMIQR and greedy batch versions of MAXV and MAXIQR (see Section 5.5), σ_n^2 needs to be also known at candidate design points $\boldsymbol{\theta}^*$. Bootstrap cannot be used because the simulated summaries are only available for training data. We take a pragmatic approach and set $\sigma_n = 10^{-2}$ at the candidate design points as if the future evaluations were almost exact. This simplification effectively reduces the occurrence of (potentially redundant) simulations at nearby points to encourage exploration. Alternatively, one could use another GP to model the bootstrapped variances or their logarithms and use the GP mean function as a point estimate for the function σ_n^2 as in Ankenman et al. (2010).

5.5 Alternative heuristic designs strategies

Here we present some heuristic alternative design strategies. These are empirically compared to the more principled EIV and IMIQR strategies in Section 6. We first focus on sequential designs where $b = 1$.

MAXIQR: A natural and simple approach is to evaluate where the current variance, IQR or some other suitable (local) measure of uncertainty is maximised. Such strategies are in some contexts called “uncertainty sampling”. The advantage over EIV and IMIQR is cheaper computation because the effect of the candidate design point to the whole

posterior is not acknowledged. Using IQR produces the design strategy

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta} \in \Theta} \pi(\boldsymbol{\theta}) e^{m_t(\boldsymbol{\theta})} \sinh(us_t(\boldsymbol{\theta})), \quad (38)$$

which we abbreviate as MAXIQR because it evaluates at the maximiser of IQR. Taking the logarithm of (38), the MAXIQR strategy can be equivalently written as

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta} \in \Theta} \left(\log \pi(\boldsymbol{\theta}) + m_t(\boldsymbol{\theta}) + us_t(\boldsymbol{\theta}) + \log(1 - e^{-2us_t(\boldsymbol{\theta})}) \right), \quad (39)$$

which shows a tradeoff between evaluating where the log-posterior is presumed to be large (the first two terms in (39)) and unexplored regions where the GP variance is large (the last two terms). This formula also shows an interesting connection to the upper confidence bound (UCB) criterion commonly used in BO, see e.g. Srinivas et al. (2010); Shahriari et al. (2015). The UCB acquisition function is $\text{UCB}(\boldsymbol{\theta}) = m_t(\boldsymbol{\theta}) + \beta_t s_t(\boldsymbol{\theta})$, where β_t is a tradeoff parameter, here automatically chosen to be $\beta_t = \Phi^{-1}(p_u)$. Compared to the standard UCB, there is, however, an extra term in (39) which further penalises regions having small variance s_t^2 . If the variance s_t^2 is large everywhere and/or if p_u is an extreme quantile, then the last term in (39) is approximately zero and the MAXIQR design criterion approximately equals UCB.

MAXV: When the variance is used instead of IQR, we obtain a strategy

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta} \in \Theta} \pi^2(\boldsymbol{\theta}) e^{2m_t(\boldsymbol{\theta}) + s_t^2(\boldsymbol{\theta})} \left(e^{s_t^2(\boldsymbol{\theta})} - 1 \right). \quad (40)$$

This strategy is abbreviated as MAXV which, in fact, is used by Gunter et al. (2014); Kandasamy et al. (2015) in the noiseless case, and it is called “exponentiated variance” by Kandasamy et al. (2015). Taking logarithm of (40) shows that this design also features a tradeoff between large posterior and large variance, similarly to MAXIQR.

Since these two strategies are not derived from Bayesian decision theory, it is not immediately clear how one should parallelise these inherently sequential strategies. However, it seems reasonable to use the fact the $s_t^2(\boldsymbol{\theta})$ is always reduced near the pending evaluation locations. Motivated by this and related BO techniques in Ginsbourger et al. (2010); Snoek et al. (2012); Desautels et al. (2014), we compute the median value of the design criterion with respect to the posterior predictive distribution of the pending simulations. The next locations are chosen iteratively such that, for MAXIQR, the first point $\boldsymbol{\theta}_1^*$ in the batch is chosen using (38) and the rest $\boldsymbol{\theta}_{2:b}$ by iteratively solving

$$\boldsymbol{\theta}_r^* = \arg \max_{\boldsymbol{\theta} \in \Theta} \pi(\boldsymbol{\theta}) e^{m_t(\boldsymbol{\theta})} \sinh(us_{t+r-1}(\boldsymbol{\theta}; \boldsymbol{\theta}_{1:r-1}^*)), \quad r = 2, 3, \dots, b. \quad (41)$$

MAXV is parallelised similarly but using the expected value instead of the median.

Finally, we provide some intuition to (41) and show a connection to the local penalisation method used to parallelise sequential BO designs by Gonzalez et al. (2016). Suppose we are selecting the r th point of a batch where $2 \leq r \leq b$. Comparison of (38) and (41) shows that (41) equals the original design criterion in (38) multiplied

by a weight function $\omega(\boldsymbol{\theta}; \boldsymbol{\theta}_{1:r-1}^*) = \sinh(us_{t+r-1}(\boldsymbol{\theta}; \boldsymbol{\theta}_{1:r-1}^*)) / \sinh(us_t(\boldsymbol{\theta}))$. It is easy to see that $\omega(\boldsymbol{\theta}; \boldsymbol{\theta}_{1:r-1}^*) \in [0, 1]$. This shows that when we take the median over the log-likelihood evaluation at the pending points $\boldsymbol{\theta}_{1:r-1}^*$, we are implicitly making the original acquisition function smaller around the pending points and, consequently, penalising additional evaluations there. This resembles the heuristic method by Gonzalez et al. (2016), who proposed to multiply the non-negative acquisition function, such as the objective of (38), with $\prod_{1 \leq j < r} \varphi(\boldsymbol{\theta}; \boldsymbol{\theta}_j^*)$, where $\varphi(\boldsymbol{\theta}; \boldsymbol{\theta}_j^*)$ are local penalising functions around the pending evaluation locations $\boldsymbol{\theta}_j^*$, when selecting the r th point $\boldsymbol{\theta}_r^*$ of the current batch. However, one difference between these approaches is that our weight function ω takes the interactions between the pending points into account and it cannot be factorised as $\omega(\boldsymbol{\theta}, \boldsymbol{\theta}_{1:r-1}^*) = \prod_{1 \leq j < r} \varphi(\boldsymbol{\theta}; \boldsymbol{\theta}_j^*)$. Also, our weight function is not a tuning parameter but follows automatically from our analysis.

6 Experiments

We empirically investigate the performance of the proposed algorithm with different design strategies developed in Section 5. We compare the sequential, batch, and greedy batch strategies based on EIV and IMIQR to sequential and greedy versions of MAXV (which is essentially the same as the BAPE method by Kandasamy et al. 2015) and MAXIQR. As a simple baseline we also sample design points from the prior (always uniform) and this method is abbreviated as RAND.

We report the results as figures whose y-axis shows the accuracy between the estimated and the ground truth posterior using total variation distance (TV). TV between pdfs π_1 and π_2 is defined as $\text{TV}(\pi_1, \pi_2) = 1/2 \int_{\Theta} |\pi_1(\boldsymbol{\theta}) - \pi_2(\boldsymbol{\theta})| d\boldsymbol{\theta}$ and is computed using numerical integration in 2D. In higher dimensional cases we compute the average TV between the marginal posterior densities using MCMC samples. The marginal median estimator in (17) is used to obtain the point estimate for the posterior pdf. The x-axis represents the iteration i of the Algorithm 1 (unless explicitly stated otherwise) which serves as a proxy to the total wall-time when the noisy likelihood evaluations are assumed to dominate the total computational cost. We use a fixed simulation budget so that the batch-sequential methods terminate earlier than the sequential ones because they spend the evaluation budget b times faster due to the parallel computation.

We consider two sets of experiments: toy models where noisy log-likelihood evaluations are directly evaluated (Section 6.1) and real-world simulator-based statistical models where SL is used to obtain noisy log-likelihood evaluations using N repeated simulations at each proposed parameter (Section 6.2). Although in the SL case it is often possible to adjust N adaptively, we use $N = 100$ (unless explicitly stated otherwise) for simplicity. In the supplementary material D.4 we show that our batch methods are beneficial for SL even though in principle it would be possible to directly parallelise the N simulations themselves. For example, when 1,000 computer cores are available for the simulations (e.g. in a high performance computing cluster), it is beneficial to use the batch strategies (e.g. $N = 100$ and $b = 10$) instead parallelising the evaluations at a single location ($N = 1000$ and $b = 1$). More elaborate analysis of resource allocation

is left for future work. A MATLAB implementation of our algorithms is available at <https://github.com/mjarvenpaa/parallel-GP-SL>.

6.1 Noisy toy model likelihoods

We first define three 2D densities with different characteristics: a simple Gaussian density called ‘Simple’, a banana-shaped density ‘Banana’ and a bimodal density ‘Bimodal’. We then construct three 6D densities so that their 2D blocks are independent and have the corresponding 2D densities as their 2D marginals. Detailed specification and illustrations can be found in supplementary material D.2. We use the same names for the 6D densities as for the corresponding 2D ones (except that ‘Bimodal’ is called ‘Multimodal’ in 6D because it has $2^3 = 8$ modes). The independence assumption is not taken into account in the GP model to make the inference problem more challenging. For simplicity, $\sigma_n(\boldsymbol{\theta})$ is assumed constant i.e. it does not depend on the magnitude of the log-likelihood and its value is obtained using MAP estimation together with other GP hyperparameters $\boldsymbol{\phi}$ at each iteration i in Algorithm 1. As an initial design in 6D we generate $b_0 = 20$ parameters (‘Simple’) or $b_0 = 50$ (‘Banana’ and ‘Multimodal’) from uniform priors. In the 2D case we always use $b_0 = 10$. We use a fixed total budget of $t = 620$ noisy log-likelihood evaluations (‘Simple’) or $t = 650$ (‘Banana’ and ‘Multimodal’) for both sequential and batch methods in 6D.

The results with different sequential and greedy batch-sequential strategies in 6D case with batch size $b = 5$ are shown in Figure 3. Good posterior approximations for the Simple example are obtained earlier than for the two other models. This is a consequence of the quadratic terms in the GP prior mean function and the exact Gaussian shape of the posterior. However, more complicated posteriors are also estimated accurately although more iterations are needed to obtain reasonable approximations. The IMIQR method works clearly the best outperforming EIV and the heuristic MAXV and MAXIQR methods which either need more iterations to obtain good approximation or fail completely to reach good results. The uniform design RAND works adequately in the Simple and Banana models but often produces poor estimates for the Multimodal case. Unsurprisingly, its performance is also poor in the real-world scenarios in Section 6.2.

The batch-sequential strategies improve the convergence speed as compared to the corresponding sequential strategies in all cases of Figure 3. In particular, the greedy batch versions of MAXV and MAXIQR even outperform the corresponding sequential methods. The greedy batch strategy in these cases encourages exploration as compared to the corresponding sequential strategy and this effect counterbalances the exploitative nature of MAXV and MAXIQR. In supplementary material D.3 we compare the joint and greedy batch strategies in 2D case where joint maximisation is still feasible. Their difference is found small for IMIQR and small or moderate for EIV suggesting that the greedy strategies are in practice nearly optimal.

Figure 4 and further examples in the supplementary material D.3 illustrate the design points and estimated posteriors for various design strategies in 2D case. An important observation is that MAXV and IMIQR are exploitative i.e. they produce




Figure 3: Results for the 6D toy densities. The lines show the median TV over 50 repeated simulations. Note that x-axis is on log-scale and the maximum number of iterations for the sequential methods is $i = 600$ and for batch methods $i = 120$.




Figure 4: The design locations for IMIQR are shown after 90 noisy log-likelihood evaluations of the 2D Banana example with noise level $\sigma_n = 1$. The black crosses show the $b_0 = 10$ initial evaluations, black dots show obtained design points except for the last two batches, and the red squares and diamonds show the last two batches.

points near the mode of the posterior where the local measure of uncertainty they use tend to be highest. Also, in general, the sequential and batch methods produce similar designs. However, greedy MAXIQR generates more points on the boundary than the corresponding sequential strategy and the joint IMIQR produces slightly more diverse design points as the sequential and greedy batch IMIQR. In all cases, IMIQR avoids redundant evaluations on the boundaries.

We investigate the effect of batch size b in the greedy batch MAXIQR and MAXIQR algorithms in Figure 5. In general, the convergence speed of both methods scales well as a function of b and $b = 10$ already yields useful improvements. However, increasing b over 40 would improve the results only slightly. Greedy batch IMIQR works overall better than batch MAXIQR. The variability in the posterior approximations produced by IMIQR is small in all cases unlike for MAXIQR which occasionally produced poor approximations (not shown for clarity). In the supplementary material D.3 we compare EIV and IMIQR in 2D. These results show that the greedy IMIQR batch-sequential




Figure 5: Results with greedy batch strategies and varying batch size b for the 6D toy models. For each method, the median TV computed over 50 repeated simulations is shown. The x-axis is truncated after $i = 400$ iterations to ease visualisation.

strategy outperforms the corresponding EIV strategy although their difference is small in the corresponding sequential cases. This suggests that, even when σ_n is small so that the variance in EIV serves as a reasonable measure of uncertainty and the sequential EIV works similarly to IMIQR, the greedy batch median-based IMIQR design strategy better mimics the sequential decisions than EIV.

6.2 Simulation models

We perform experiments with three benchmark problems used previously in the ABC literature. Two of these are shown here and the third one in the supplementary material D.5. While the proposed methodology is particularly useful for expensive simulation models, we however consider only relatively cheap models as this allows to repeat the computations many times with different realisations of randomness to assess the variability and robustness, and to conduct accurate comparisons to reasonable ground truth posteriors. Nevertheless, these experiments serve as examples of challenging real-world inference scenarios where the GP and SL modelling assumptions do not hold exactly. In each problem, we set the unknown parameter of the simulation model to a value used previously in the literature and generated one data set from the simulation model using this “true” parameter. The posterior used as the ground truth was computed using SL-MCMC. Multiple chains each with length 10^6 were used to ensure that the variability due to Monte Carlo error was small.

Ricker model

We first consider the Ricker model presented in Wood (2010). In this model N_t denotes the number of individuals in a population at time t which evolves according to the discrete time stochastic process $N_{t+1} = rN_t \exp(-N_t + \varepsilon_t)$, for $t = 1, \dots, T$, where $\varepsilon_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2)$. The initial population size is $N_0 = 1$. It is assumed that only a noisy




Figure 6: Results for the Ricker model. (a) The median TV and 90% variability interval over 100 repeated runs of the algorithms. (b-d) Estimated posterior marginal densities (black) shown for five typical runs of the algorithm with the greedy batch IMIQR strategy. The ground truth computed with SL-MCMC (red) and the true parameter value (red diamond) are also shown for comparison.

measurement x_t of the population size N_t at each time point is available with the Poisson observation model $x_t | N_t, \phi \sim \text{Poi}(\phi N_t)$. Given data $\mathbf{x} = (x_t)_{t=1}^T$, the goal is to infer the three parameters $\boldsymbol{\theta} = (\log(r), \phi, \sigma_\epsilon)$. We use the uniform prior $(\log(r), \phi, \sigma_\epsilon) \sim \mathcal{U}([3, 5] \times [4, 20] \times [0, 0.8])$. The same 13 summary statistics as in Wood (2010); Gutmann and Corander (2016); Price et al. (2018) are used to compute log-SL evaluations. The number of repeated simulations is fixed to $N = 100$. The “true” parameter to be estimated is $\boldsymbol{\theta}_{\text{true}} = (3.8, 10, 0.3)$ and it is used to generate the observed data with length $T = 50$. The initial training data size is $b_0 = 30$ and the additional budget of simulations is 600 so that the total budget is 630 SL evaluations corresponding 63000 simulations. The integrals of EIV and IMIQR are approximated using IS and σ_n^2 is estimated using bootstrap as described in Section 5.4.

Figure 6 shows the results. We see that the EIV and MAXV strategies perform poorly. These strategies tend to evaluate where the variance of the posterior is high, although as discussed in Section 5, these do not necessarily correspond to the regions with non-negligible likelihood. In fact, the magnitude of the log-likelihood and its noise variance σ_n^2 grow fast near the boundaries of the parameter space where the chaotic nature of the model also makes the log-likelihood surface irregular which further causes difficulties with GP modelling. The IMIQR method again produces the best posterior approximations which are comparable to the true SL posterior. Some examples are shown in Figure 6b-d. Also, MAXIQR method works well on average but it produces less coherent results than IMIQR which is likely the result of its exploitative nature. In addition, unexpectedly, the greedy MAXIQR method performs poorly. The probable reason is that the batch evaluations become too diverse having many evaluations in the boundary which leads to poor GP fitting and subsequently poor future designs. However, the robust batch-sequential IMIQR method with $b = 5$ works as expected producing useful improvement to the convergence speed as compared to sequential IMIQR.

g-and-k model

We consider the g-and-k distribution as in Price et al. (2018). The g-and-k model is a flexible probability distribution defined via its quantile function

$$Q(\Phi^{-1}(p); \boldsymbol{\theta}) = a + b \left(1 + c \frac{1 - \exp(-g\Phi^{-1}(p))}{1 + \exp(-g\Phi^{-1}(p))} \right) (1 + (\Phi^{-1}(p))^2)^k \Phi^{-1}(p), \quad (42)$$

where a, b, c, g and k are parameters and $p \in [0, 1]$ is a quantile. We fix $c = 0.8$ and estimate the parameters $\boldsymbol{\theta} = (a, b, g, k)$ using a uniform prior $\pi(\boldsymbol{\theta}) = \mathcal{U}([2.5, 3.5] \times [0.5, 1.5] \times [1.5, 2.5] \times [0.3, 0.7])$. We use the same four summary statistics as Price et al. (2018) who fitted an auxiliary model, skew t -distribution, to the set of samples generated from (42) using maximum likelihood, and took the resulting skew t score vector at the ML estimate as the summary statistic. Although there are only 4 summary statistics, we again use $N = 100$. We use the same settings as for the Ricker model except that the initial design is increased to $b_0 = 40$ so that the total budget is 640 SL evaluations. The true value of the parameter is chosen to be $\boldsymbol{\theta}_{\text{true}} = (3, 1, 2, 0.5)$.

Overall, the results in Figure 7 are similar to those of the Ricker model. However, the larger parameter space slows down the convergence speeds initially as expected, as compared to the Ricker model. Low dimension of the summary statistic and the moderately large value $N = 100$ cause the log-likelihood evaluations to be quite accurate near the modal area of the likelihood ($\sigma_n(\boldsymbol{\theta}_{\text{true}}) \approx 0.15$) and we expect that smaller N might be already enough. However, using $N = 100$ ensures accurate variance estimates using




Figure 7: (a) Results for the g-and-k model. (a) Median TV and 90% variability interval over 100 repeated runs. (b-e) Some estimated posterior marginal densities illustrated as in Figure 6.

the bootstrap. While MAXIQR strategy works almost as well as IMIQR on average, it completely fails in some individual repeated experiments producing long variability intervals in Figure 7 leaving IMIQR as the only successful method.

Effect of batch size

As the last experiment, we investigate the improvements brought by the batch-sequential IMIQR strategy in the case of real-world simulation models. We use the Ricker and g-and-k models from the previous subsections. The experiment details are the same except that we consider only IMIQR strategy with several batch sizes $b \in \{2, 5, 10, 20, 30\}$. The results in Figure 8 show that, on average, the greedy batch-sequential IMIQR with batch sizes up to 30 produces as good approximations as the corresponding sequential strategy. The convergence speed is also improved almost linearly.

However, the variability in the quality of the estimated posteriors increases with larger batch sizes when the total budget of simulations is kept fixed. While most of the repeated runs of the algorithm have converged to excellent approximations in all cases as seen in Figure 8, there were some individual runs where the algorithm did not yet converge when the budget was used. The posterior estimate at the final iteration is often quite poor in these cases. Most of these happen with Ricker model when $b \geq 20$ and with g-and-k model when $b = 30$. However, this behaviour is not surprising: When b is large, the complete batch is constructed using the same limited information which




Figure 8: Results of the greedy batch strategy IMIQR with various batch sizes b for Ricker and g-and-k models. Top row shows the median TV over 100 repeated experiments and the bottom row shows the corresponding 90% quantile.

necessarily produces occasional poor batches providing little information. Furthermore, the importance density in (37) is likely to get worse when the batch size is increased and cause the last points in the batch to be less useful. It is thus inevitable that the batch size should not be chosen too large. Nevertheless, it is seen that batch size $b = 10$ already produces substantial gains (especially considering that further parallelisation is often possible with respect to N) and produces consistently accurate posterior approximations.

7 Discussion and conclusions

If only a limited number of noisy log-likelihood evaluations can be computed, standard techniques such as MCMC become difficult to use for Bayesian inference. To tackle the problem, we constructed a hierarchical GP surrogate model for the noisy log-likelihood and discussed properties of the resulting estimators of the (unnormalised) posterior. We developed two batch-sequential strategies (EIV and IMIQR) based on Bayesian decision theory, to (semi-)optimally select the next evaluation locations and to parallelise the costly simulations. We also considered heuristic design strategies (MAXV and MAXIQR). We provided some theoretical analysis: We derived an approximate bound for the greedy optimisation of the batch IMIQR method using the concept of weak submodularity, showed a connection between the UCB (a common BO method) and the MAXIQR strategy, and between batch MAXIQR and the local penalisation method by Gonzalez et al. (2016). The proposed methods were investigated experimentally.

The IMIQR strategy was found to be robust both to violations of the GP surrogate model assumptions and to the heavy-tails of the resulting distributions. Unlike the other design strategies, it consistently produced posterior approximations comparable to the ground truth. Greedy batch-sequential IMIQR strategy was found to be highly useful to parallelise the potentially expensive simulations. In our experiments it produced substantial, sometimes even linear, speed improvements for batch sizes $b = 20$. We thus recommend the IMIQR strategy. In general we were able to obtain useful posterior approximations with 10,000 to 20,000 simulations that can be easily parallelised. This is considerably less than e.g. using (pseudo-marginal) MCMC requiring typically at least tens of thousands of iterations corresponding to millions of simulations, careful convergence assessment and tuning of the proposal density. Another important observation was that the heuristic strategies that evaluate where the current uncertainty is highest, despite their small computational cost and good performance in earlier studies with deterministic evaluations, worked poorly with noisy log-likelihood evaluations.

Similarly to other GP surrogate techniques such as BO, fitting the GP and finding the next evaluation locations by optimising the design criterion is however not free. Our unoptimised MATLAB implementations of MAXV and MAXIQR are fast, but optimisation of the EIV and IMIQR design criteria takes a couple of seconds in 2D and around 20 to 80 seconds per parameter in 3D and 4D. This means that the proposed algorithm is useful when the simulation time is several seconds or more, which is however true with many real-world simulation models. Furthermore, the quality of the posterior

approximation also depends on the choice of the surrogate model. We used the same GP model in all of our experiments with no problem-specific tuning, which already produced good results. However, some problems would certainly benefit from further adjustment and incorporation of domain knowledge. For example, if the likelihood is expected to be flat, a GP prior with a constant mean function might be appropriate.

In this work, similarly to Rasmussen (2003); Wilkinson (2014); Kandasamy et al. (2015); Gutmann and Corander (2016); Drovandi et al. (2018), we built our surrogate GP model for the log-likelihood. An alternative way would be to model the summary statistics. Meeds and Welling (2014) used such an approach but they assumed that the summary statistics are independent. However, modelling the scalar-valued log-likelihood is simpler and our approach also applies as such to non-ABC scenarios with exact (but potentially expensive) log-likelihood evaluations as in Osborne et al. (2012); Kandasamy et al. (2015); Wang and Li (2018); Acerbi (2018). Gutmann and Corander (2016); Järvenpää et al. (2018, 2019) modelled the discrepancy between simulated and observed data with a GP and obtained reasonable posterior approximations with only a few hundred model simulations. Here we need N repeated simulations just to compute the log-likelihood for a single parameter value, which can be seen as the price of not having to specify an explicit discrepancy measure and the ABC tolerance.

We see several avenues for future research. The consistency and convergence rates of our algorithms could be investigated theoretically. Some work towards that direction has been done by Bect et al. (2019); Stuart and Teckentrup (2018). Adaptive control of the number of repeated simulations N could likely be used to further reduce the number of simulations required, possibly as in Picheny et al. (2013). Some simulation models may behave unexpectedly near the boundaries of the parameter space violating GP model assumptions as we saw with the Ricker model in Section 6. Similarly, situations where the prior is significantly more diffuse than the posterior may be unsuitable for our approach that relies on a global GP surrogate. Consequently, it would be useful to learn adaptively not only where to evaluate next but also which parameter regions to rule out completely. This could be done as in Wilkinson (2014) or possibly by adapting ideas from the constrained BO literature (Gardner et al., 2014; Sui et al., 2015).

Supplementary Material

Supplementary material of “Parallel Gaussian process surrogate Bayesian inference with noisy likelihood evaluations”

(DOI: [10.1214/20-BA1200SUPP](https://doi.org/10.1214/20-BA1200SUPP); .pdf). The supplementary material contains proofs, implementation details and additional experimental results.

References

- Acerbi, L. (2018). “Variational Bayesian Monte Carlo.” In *Advances in Neural Information Processing Systems 31*, 8223–8233. [148](#), [158](#), [172](#)
- An, Z., Nott, D. J., and Drovandi, C. (2019a). “Robust Bayesian synthetic likelihood via

- a semi-parametric approach.” *Statistics and Computing*. MR4065218. doi: <https://doi.org/10.1007/s11222-019-09904-x>. 150
- An, Z., South, L. F., Nott, D. J., and Drovandi, C. C. (2019b). “Accelerating Bayesian Synthetic Likelihood with the Graphical Lasso.” *Journal of Computational and Graphical Statistics*, 28(2): 471–475. MR3974895. doi: <https://doi.org/10.1080/10618600.2018.1537928>. 150
- Ankenman, B., Nelson, B. L., and Staum, J. (2010). “Stochastic Kriging for Simulation Metamodeling.” *Operations Research*, 58(2): 371–382. MR2674803. doi: <https://doi.org/10.1287/opre.1090.0754>. 162
- Azimi, J., Alan, F., and Fern, X. Z. (2010). “Batch Bayesian Optimization via Simulation Matching.” In *Advances in Neural Information Processing Systems 23*, 109–117. 148
- Bach, F. (2013). *Learning with Submodular Functions: A Convex Optimization Perspective*. Hanover, MA, USA: Now Publishers Inc. 159
- Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). “Approximate Bayesian computation in population genetics.” *Genetics*, 162(4): 2025–2035. 147
- Bect, J., Bachoc, F., and Ginsbourger, D. (2019). “A supermartingale approach to Gaussian process based sequential design of experiments.” *Bernoulli*, 25(4A): 2883–2919. MR4003568. doi: <https://doi.org/10.3150/18-BEJ1074>. 172
- Bect, J., Ginsbourger, D., Li, L., Picheny, V., and Vazquez, E. (2012). “Sequential design of computer experiments for the estimation of a probability of failure.” *Statistics and Computing*, 22(3): 773–793. MR2909621. doi: <https://doi.org/10.1007/s11222-011-9241-4>. 149, 157
- Briol, F.-X., Oates, C. J., Girolami, M., Osborne, M. A., and Sejdinovic, D. (2019). “Probabilistic Integration: A Role in Statistical Computation?” *Statistical Science*, 34(1): 1–22. MR3938958. doi: <https://doi.org/10.1214/18-STS660>. 152
- Brochu, E., Cora, V. M., and de Freitas, N. (2010). “A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning.” Available at: <https://arxiv.org/abs/1012.2599>. 148
- Chai, H. R. and Garnett, R. (2019). “Improving Quadrature for Constrained Integrands.” In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2751–2759. 148
- Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., and Richet, Y. (2014). “Fast Parallel Kriging-Based Stepwise Uncertainty Reduction With Application to the Identification of an Excursion Set.” *Technometrics*, 56(4): 455–465. MR3290615. doi: <https://doi.org/10.1080/00401706.2013.860918>. 149, 161, 162
- Cockayne, J., Oates, C., Sullivan, T., and Girolami, M. (2019). “Bayesian Probabilistic Numerical Methods.” *SIAM Review*, 61(4): 756–789. MR4027836. doi: <https://doi.org/10.1137/17M1139357>. 152

- Contal, E., Buffoni, D., Robicquet, A., and Vayatis, N. (2013). “Parallel Gaussian process optimization with upper confidence bound and pure exploration.” In *Lecture Notes in Computer Science*. 148
- Desautels, T., Krause, A., and Burdick, J. W. (2014). “Parallelizing Exploration-Exploitation Tradeoffs in Gaussian Process Bandit Optimization.” *Journal of Machine Learning Research*, 15: 4053–4103. MR3317214. 148, 163
- Drovandi, C. C., Moores, M. T., and Boys, R. J. (2018). “Accelerating pseudo-marginal MCMC using Gaussian processes.” *Computational Statistics & Data Analysis*, 118: 1–17. MR3715260. doi: <https://doi.org/10.1016/j.csda.2017.09.002>. 148, 151, 172
- Frazier, D. T., Nott, D. J., Drovandi, C., and Kohn, R. (2019). “Bayesian inference using synthetic likelihood: asymptotics and adjustments.” Available at: <https://arxiv.org/abs/1902.04827>. 150
- Gardner, J., Kusner, M., Zhixiang, Weinberger, K., and Cunningham, J. (2014). “Bayesian Optimization with Inequality Constraints.” In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, 937–945. 172
- Ginsbourger, D., Le Riche, R., and Carraro, L. (2010). *Kriging Is Well-Suited to Parallelize Optimization*, 131–162. Berlin, Heidelberg: Springer Berlin Heidelberg. 148, 158, 159, 163
- Gonzalez, J., Dai, Z., Lawrence, N. D., and Hennig, P. (2016). “Batch Bayesian Optimization via Local Penalization.” In *International Conference on Artificial Intelligence and Statistics*, 1, 648–657. 148, 163, 164, 171
- González, J., Osborne, M., and Lawrence, N. D. (2016). “GLASSES: Relieving The Myopia of Bayesian Optimisation.” In *Proceedings of the Nineteenth International Workshop on Artificial Intelligence and Statistics*. 157
- Gunter, T., Osborne, M. A., Garnett, R., Hennig, P., and Roberts, S. J. (2014). “Sampling for Inference in Probabilistic Models with Fast Bayesian Quadrature.” In *Advances in Neural Information Processing Systems 27*, 2789–2797. 148, 158, 163
- Gutmann, M. U. and Corander, J. (2016). “Bayesian optimization for likelihood-free inference of simulator-based statistical models.” *Journal of Machine Learning Research*, 17(125): 1–47. MR3555016. 148, 149, 151, 168, 172
- Hennig, P., Osborne, M. A., and Girolami, M. (2015). “Probabilistic numerics and uncertainty in computations.” *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 471(2179): 20150142. MR3378744. doi: <https://doi.org/10.1098/rspa.2015.0142>. 148, 152
- Hennig, P. and Schuler, C. J. (2012). “Entropy Search for Information-Efficient Global Optimization.” *Journal of Machine Learning Research*, 13(1999): 1809–1837. MR2956343. 161
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. (2014). “Predictive

- Entropy Search for Efficient Global Optimization of Black-box Functions.” *Advances in Neural Information Processing Systems 28*, 1–9. 161
- Jabot, F., Lagarrigues, G., Courbaud, B., and Dumoulin, N. (2014). “A comparison of emulation methods for Approximate Bayesian Computation.” Available at: <http://arxiv.org/abs/1412.7560>. 148
- Järvenpää, M., Gutmann, M. U., Pleska, A., Vehtari, A., and Marttinen, P. (2019). “Efficient Acquisition Rules for Model-Based Approximate Bayesian Computation.” *Bayesian Analysis*, 14(2): 595–622. MR3934099. doi: <https://doi.org/10.1214/18-BA1121>. 148, 149, 152, 154, 157, 158, 162, 172
- Järvenpää, M., Gutmann, M. U., Vehtari, A., and Marttinen, P. (2018). “Gaussian process modelling in approximate Bayesian computation to estimate horizontal gene transfer in bacteria.” *The Annals of Applied Statistics*, 12(4): 2228–2251. MR3875699. doi: <https://doi.org/10.1214/18-AOAS1150>. 147, 172
- Järvenpää, M., Gutmann, M. U., Vehtari, A., and Marttinen, P. (2020). “Parallel Gaussian Process Surrogate Bayesian Inference with Noisy Likelihood Evaluations – Supplementary Material.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/20-BA1200SUPP>. 149
- Kandasamy, K., Schneider, J., and Póczos, B. (2015). “Bayesian active learning for posterior estimation.” In *International Joint Conference on Artificial Intelligence*, 3605–3611. 148, 149, 158, 163, 164, 172
- Karvonen, T., Oates, C. J., and Särkkä, S. (2018). “A Bayes-Sard Cubature Method.” In *Advances in Neural Information Processing Systems 31*, 5886–5897. MR4026668. doi: <https://doi.org/10.1007/s11222-019-09896-8>. 148
- Kennedy, M. C. and O’Hagan, A. (2001). “Bayesian calibration of computer models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3): 425–464. MR1858398. doi: <https://doi.org/10.1111/1467-9868.00294>. 147
- Krause, A. and Cevher, V. (2010). “Submodular Dictionary Selection for Sparse Representation.” In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 567–574. 160
- Krause, A., Singh, A., and Guestrin, C. (2008). “Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies.” *Journal of Machine Learning Research*, 9: 235–284. 159, 160
- Lintusaari, J., Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. (2017). “Fundamentals and Recent Developments in Approximate Bayesian Computation.” *Systematic biology*, 66(1): e66–e82. 147, 149
- Lyu, X., Binois, M., and Ludkovski, M. (2018). “Evaluating Gaussian Process Meta-models and Sequential Designs for Noisy Level Set Estimation.” Available at: <http://arxiv.org/abs/1807.06712>. 149, 154
- Marin, J. M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). “Approximate Bayesian

- computational methods." *Statistics and Computing*, 22(6): 1167–1180. MR2992292. doi: <https://doi.org/10.1007/s11222-011-9288-2>. 147, 149
- Marttinen, P., Gutmann, M. U., Croucher, N. J., Hanage, W. P., and Corander, J. (2015). "Recombination produces coherent bacterial species clusters in both core and accessory genomes." *Microbial Genomics*, 1(5). 147
- Meeds, E. and Welling, M. (2014). "GPS-ABC: Gaussian Process Surrogate Approximate Bayesian Computation." In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*. 148, 172
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). "An analysis of approximations for maximizing submodular set functions—I." *Mathematical Programming*, 14(1): 265–294. MR0503866. doi: <https://doi.org/10.1007/BF01588971>. 159
- O'Hagan, A. (1991). "Bayes-Hermite quadrature." *Journal of Statistical Planning and Inference*. MR1144171. doi: [https://doi.org/10.1016/0378-3758\(91\)90002-V](https://doi.org/10.1016/0378-3758(91)90002-V). 148
- O'Hagan, A. and Kingman, J. F. C. (1978). "Curve Fitting and Optimal Design for Prediction." *Journal of the Royal Statistical Society. Series B (Methodological)*, 40(1): 1–42. MR0512140. 151
- Osborne, M. A., Duvenaud, D., Garnett, R., Rasmussen, C. E., Roberts, S. J., and Ghahramani, Z. (2012). "Active Learning of Model Evidence Using Bayesian Quadrature." *Advances in Neural Information Processing Systems 26*, 1–9. 148, 172
- Picheny, V., Ginsbourger, D., Richet, Y., and Caplin, G. (2013). "Quantile-Based Optimization of Noisy Computer Experiments With Tunable Precision." *Technometrics*, 55(1): 2–13. MR3038476. doi: <https://doi.org/10.1080/00401706.2012.707580>. 172
- Price, L. F., Drovandi, C. C., Lee, A., and Nott, D. J. (2018). "Bayesian Synthetic Likelihood." *Journal of Computational and Graphical Statistics*, 27(1): 1–11. MR3788296. doi: <https://doi.org/10.1080/10618600.2017.1302882>. 148, 149, 150, 168, 169
- Rasmussen, C. E. (2003). "Gaussian Processes to Speed up Hybrid Monte Carlo for Expensive Bayesian Integrals." *Bayesian Statistics 7*, 651–659. MR2003529. 148, 172
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. The MIT Press. MR2514435. 147, 151
- Riihimäki, J. and Vehtari, A. (2014). "Laplace Approximation for Logistic Gaussian Process Density Estimation and Regression." *Bayesian Analysis*, 9(2): 425–448. MR3217002. doi: <https://doi.org/10.1214/14-BA872>. 151
- Robert, C. P. (2007). *The Bayesian Choice*. New York: Springer, second edition. MR2723361. 152
- Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. New York: Springer, second edition. MR2080278. doi: <https://doi.org/10.1007/978-1-4757-4145-2>. 147

- Shah, A. and Ghahramani, Z. (2015). “Parallel Predictive Entropy Search for Batch Global Optimization of Expensive Objective Functions.” In *Advances in Neural Information Processing Systems 28*, 12. 148
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. (2015). “Taking the human out of the loop: A review of Bayesian optimization.” *Proceedings of the IEEE*, 104(1). 148, 163
- Sinsbeck, M. and Nowak, W. (2017). “Sequential Design of Computer Experiments for the Solution of Bayesian Inverse Problems.” *SIAM/ASA Journal on Uncertainty Quantification*, 5(1): 640–664. MR3679325. doi: <https://doi.org/10.1137/15M1047659>. 148, 152, 157, 158
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). “Practical Bayesian optimization of machine learning algorithms.” In *Advances in Neural Information Processing Systems 25*, 1–9. 148, 159, 163
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2010). “Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design.” In *Proceedings of the 27th International Conference on Machine Learning*, 1015–1022. MR2952544. doi: <https://doi.org/10.1109/TIT.2011.2182033>. 163
- Stuart, A. M. and Teckentrup, A. L. (2018). “Posterior consistency for Gaussian process approximations of Bayesian posterior distributions.” *Mathematics for Computing*, 87: 721–753. MR3739215. doi: <https://doi.org/10.1090/mcom/3244>. 154, 172
- Sui, Y., Gotovos, A., Burdick, J., and Krause, A. (2015). “Safe Exploration for Optimization with Gaussian Processes.” In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 997–1005. 172
- Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. (2018). “Likelihood-free inference by ratio estimation.” Available at: <https://arxiv.org/abs/1611.10242>. 150
- Turner, B. M. and Van Zandt, T. (2012). “A tutorial on approximate Bayesian computation.” *Journal of Mathematical Psychology*, 56(2): 69–85. MR2909506. doi: <https://doi.org/10.1016/j.jmp.2012.02.005>. 149
- Wang, H. and Li, J. (2018). “Adaptive Gaussian Process Approximation for Bayesian Inference with Expensive Likelihood Functions.” *Neural Computation*, 30(11): 3072–3094. MR3873817. doi: https://doi.org/10.1162/neco_a_01127. 148, 172
- Wilkinson, R. D. (2014). “Accelerating ABC methods using Gaussian processes.” In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*. 148, 151, 162, 172
- Wilson, J., Hutter, F., and Deisenroth, M. (2018). “Maximizing acquisition functions for Bayesian optimization.” In *Advances in Neural Information Processing Systems 31*, 9906–9917. 148, 159
- Wood, S. N. (2010). “Statistical inference for noisy nonlinear ecological dynamic systems.” *Nature*, 466: 1102–1104. 148, 149, 167, 168

- Wu, J. and Frazier, P. (2016). “The Parallel Knowledge Gradient Method for Batch Bayesian Optimization.” In *Advances in Neural Information Processing Systems 29*, 3126–3134. [MR3755126](#). 148, 161
- Yu, C. W. and Clarke, B. (2011). “Median loss decision theory.” *Journal of Statistical Planning and Inference*, 141(2): 611–623. [MR2732931](#). doi: <https://doi.org/10.1016/j.jspi.2010.08.013>. 158

Acknowledgments

The authors are grateful to an associate editor and two anonymous reviewers for their constructive feedback that helped to significantly improve this article. This work was funded by the Academy of Finland (grants no. 286607 and 294015 to PM). We acknowledge the computational resources provided by Aalto Science-IT project.