



HAL
open science

A portfolio approach to massively parallel Bayesian optimization

Mickael Binois, Nicholson Collier, Jonathan Ozik

► **To cite this version:**

Mickael Binois, Nicholson Collier, Jonathan Ozik. A portfolio approach to massively parallel Bayesian optimization. 2021. hal-03383097v1

HAL Id: hal-03383097

<https://inria.hal.science/hal-03383097v1>

Preprint submitted on 18 Oct 2021 (v1), last revised 3 Apr 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A portfolio approach to massively parallel Bayesian optimization

Mickaël Binois* Nicholson Collier^{†‡} Jonathan Ozik^{†‡}

Abstract

One way to reduce the time of conducting optimization studies is to evaluate designs in parallel rather than just one-at-a-time. For expensive-to-evaluate black-boxes, batch versions of Bayesian optimization have been proposed. They work by building a surrogate model of the black-box that can be used to select the designs to evaluate efficiently via an infill criterion. Still, with higher levels of parallelization becoming available, the strategies that work for a few tens of parallel evaluations become limiting, in particular due to the complexity of selecting more evaluations. It is even more crucial when the black-box is noisy, necessitating more evaluations as well as repeating experiments. Here we propose a scalable strategy that can keep up with massive batching natively, focused on the exploration/exploitation trade-off and a portfolio allocation. We compare the approach with related methods on deterministic and noisy functions, for mono and multi-objective optimization tasks. These experiments show similar or better performance than existing methods, while being orders of magnitude faster.

1 Introduction

Rather than increasing the computation speed, current trends are on parallelization on highly concurrent computing systems. The benefit for computer models (a.k.a. simulators) is geared towards increasing their accuracy rather than their time to solution. These simulators are prevalent in many fields, ranging from physics to biology or in engineering. Still, increasing the parallelization often comes with diminishing returns and evaluation time remains limiting. A strategy is then to conduct several evaluations at the same time, in batches, to optimize

*Corresponding author: Inria, Université Côte d’Azur, CNRS, LJAD, Sophia Antipolis, France
mickael.binois@inria.fr

[†]Argonne National Laboratory, Lemont, IL, USA

[‡]Consortium for Advanced Science and Engineering, University of Chicago, Chicago, IL, USA

(here minimize) quantities of interest. See for instance (Haftka et al., 2016) for a review of options.

For fast simulators, evolutionary algorithms are amenable to parallelization by design, see e.g., (Emmerich and Deutz, 2018) for a review. But they require a prohibitive number of evaluations for more expensive-to-evaluate simulators. For these, Bayesian optimization (BO), see e.g., (Shahriari et al., 2016; Frazier, 2018), is preferred, with its ability to carefully select the next evaluations. Typically, BO relies on a Gaussian process model of the simulator, or any black-box, by using a probabilistic surrogate model that can be exploited to efficiently perform the so-called exploration/exploitation trade-off. Exploitation refers to areas where the prediction is low (for minimization), while exploration is for areas of large predictive variance. An infill criterion, or acquisition function, balances this trade-off to select evaluations, such as the expected improvement (EI) (Mockus et al., 1978) in the efficient global optimization algorithm (Jones et al., 1998). Other options include upper confidence bound (UCB) (Srinivas et al., 2010), knowledge gradient (Frazier, 2018), and entropy based criteria (Villemonteix et al., 2009b; Hennig and Schuler, 2012; Wang and Jegelka, 2017). Noisy simulators have their own set of challenges, see e.g., (Baker et al., 2020), and raise questions about selecting the right amount of replication. While not necessary per se, repeating experiments is the best option to separate signal from noise, and is beneficial in terms of computational speed, see e.g., (Binois et al., 2019; Zhang et al., 2020).

Many batch versions of these infill criteria have been proposed, such as (Kandasamy et al., 2018; Hernández-Lobato et al., 2017) for Thompson sampling (TS). For EI, rather than just looking at its local optima as in (Sóbester et al., 2004), efficient heuristics propose to select batch points iteratively, replacing unknown values at selected points by pseudo-values (Ginsbourger et al., 2010). This was later coined as “hallucination” in the UCB version of (Desautels et al., 2014). Exact expressions have later been derived for the multi-point EI (qEI) by (Chevalier and Ginsbourger, 2013) and for its gradient as well by (Marmin

et al., 2015). A faster but approximate version of qEI was proposed by (Binois, 2015) and a stochastic approximation framework by (Wang et al., 2020). Rontsis et al. (2020) use an optimistic bound on EI for all possible distributions compatible with the same first two moments as a GP, which is amenable to reliable computation at the cost of solving a semi-definite problem, limiting the scaling up to large batches. Reducing the cost of selecting batches, in particular from modeling the joint probability distribution of the batch or within a hallucination scheme has been proposed by (Gonzalez et al., 2016). Their idea is to select batch members sequentially by penalizing proximity to the previously selected ones. Taking different criteria is an option to select different trade-offs, as by (Tran et al., 2019). This idea of a portfolio of acquisition functions is also present in (Hoffman et al., 2011), but limited to a few options. Using local models is another option to select batches efficiently, up to several hundreds in (Wang et al., 2018). The downside is a lack of coordination in the selection and the need of an *ad hoc* selection procedure. In the case of entropy or stepwise uncertainty reduction criteria, see, e.g., (Chevalier et al., 2014), batching would increase their already intensive computational burden. Another early work by (Azimi et al., 2010) attempts to match the expected sequential performance, via approximations and sampling.

Here, we also consider multi-objective optimization (MOO), motivated by a calibration problem of a large-scale agent-based model (ABM) of COVID-19 run on a HPC cluster. There the goal is to find the set of best compromise solutions, the Pareto front, since there is rarely a solution minimizing all objectives at once. See, e.g., (Hunter et al., 2017) for a review of options for black-boxes and (Emmerich et al., 2020) for a discussion on the multi-objective (MO) BO infill criteria. MO versions of multi-point algorithms have also been proposed, taking different scalarization weights (Zhang et al., 2010), or relying on an additional notion of diversity as in (Lukovic et al., 2020). Classifying Pareto points is also considered, see e.g., (Zuluaga et al., 2016).

The operations research community has seen works dealing with low signal-to-noise ratios

and heteroskedasticity, where replication is key. Generally, the idea is to first identify good points before defining the number of replicates, see, e.g., (Zhang et al., 2017; Gonzalez et al., 2020) or (Rojas-Gonzalez and Van Nieuwenhuysse, 2020) for a review on stochastic MO BO. Still, the batch aspect is missing in the identification of candidates.

There are several limitations in the works above. First, while it is generally assumed that the cost of one evaluation is sufficiently high to consider the time to select new points negligible, this may not be the case in the large batch setup. Indeed, evaluation budgets can be larger when many points can be evaluated at once, increasing the cost of fitting GPs (cubic in the number of designs, for the vanilla version). Then parallel infill criteria are more expensive to evaluate. Second, optimizing these parallel infill criteria is increasingly harder. That is, frontally optimizing the acquisition function requires being able to solve a global optimization problem in dimension (d) number of variables \times batch size (q). But even solving it approximately one design at a time becomes costly as it must be done sequentially q times. Both become cumbersome for large batches and, presumably, far from the global optimum since the acquisition function landscape is multimodal, with flat regions, and symmetry properties. Plus, as we showcase, this results in parts of batch members being less pertinent. Many works also focus on discrete search spaces, which can be inefficient as the number of variables increases. Lastly, adaptive batch sizes might be more efficient than a fixed number of parallel evaluations, see e.g., (Desautels et al., 2014). Similarly, asynchronous evaluation is another angle to exploit when the evaluation times vary, see e.g., (Gramacy and Lee, 2009; Janusevskis et al., 2012; Alvi et al., 2019).

The targeted setup is as follows: a massively parallel system, with the option to run hundreds of evaluations in parallel over several iterations (e.g., with an overall job duration as a stopping criteria). This is sometimes called the high throughput regime (Hernández-Lobato et al., 2017). Hence the time dedicated to select the batch points should be minimal (and not considered negligible), as well as amenable to parallelization (to use most of the whole system

during the job). The method we propose is to directly identify candidates realizing different exploration/exploitation trade-offs. This amounts to approximating the GP predictive mean-variance Pareto front, which is orders of magnitude faster than most existing infill criteria. In the noisy case, possibly with input-dependent variance, the variance reduction makes an additional objective to encourage replication. Then, to actually select batches, we follow the approach proposed by (Guerreiro and Fonseca, 2016) with the hypervolume Sharpe ratio indicator (HSRI) in the context of evolutionary algorithms. In the MO version, the transposition is to directly take the mean and variance of each objective.

The contributions of this work are:

- a simple portfolio allocation strategy relying on the exploration/exploitation trade-off, plus variance reduction in the noisy case. It extends directly to the multi-objective setup;
- an approach independent of the size of the batch, removing limitations of current batch criteria for large q (computational complexity, difficult optimization, low sensitivity to some designs);
- the potential for flexible batch sizes and asynchronous allocation via the portfolio approach;
- the ability to natively take into account replication.

In Section 2 we briefly present GPs, and batch and MO BO. In Section 3, the difficulties of batch BO are described before detailing the proposed method. It is then tested and compared empirically with alternatives in Section 4. A conclusion is given in Section 5.

2 Background

We consider the minimization problem: find $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^d} f(\mathbf{x})$ where \mathbb{X} is typically a hypercube, $f : \mathbb{X} \rightarrow \mathbb{R}$. Among a variety of options for surrogate modeling of the black-box f , see e.g., Shahriari et al. (2016), GPs are prevalent.

2.1 Gaussian process regression

Consider a set of $n \in \mathbb{N}^*$ designs-observations couples (\mathbf{x}_i, y_i) with $y_i = f(\mathbf{x}_i) + \varepsilon_i$, $\varepsilon_i \sim \mathcal{N}(0, r(\mathbf{x}_i))$, often obtained with a Latin hypercube sample (LHS) as design of experiments (DoE). The idea of GP regression, or kriging, is to assume that f follows a multivariate normal distribution, characterized by an arbitrary mean function $m(\mathbf{x})$ and a positive semi-definite covariance kernel function $k : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$. Unless prior information is available to specify a mean function, m is assumed to be zero for simplicity. As for k , parameterized families of covariance functions such as Gaussian or Matérn ones are preferred, whose hyperparameters (process variance, lengthscales) can be inferred in many ways, such as maximum likelihood estimation.

Conditioned on observations $\mathbf{y} := (y_1, \dots, y_n)$, zero mean GP predictions at any set of q designs in \mathbb{X} , $\mathcal{X}_q : (\mathbf{x}'_1, \dots, \mathbf{x}'_q)^\top$, are still Gaussian, $Y_n(\mathcal{X}_q) | \mathbf{y} \sim \mathcal{N}(m_n(\mathcal{X}_q), s_n^2(\mathcal{X}))$:

$$\begin{aligned} m_n(\mathcal{X}_q) &= \mathbf{k}_n(\mathcal{X}_q)^\top \mathbf{K}_n^{-1} \mathbf{y}_n \\ s_n^2(\mathcal{X}_q) &= k(\mathcal{X}_q, \mathcal{X}_q) - \mathbf{k}_n(\mathcal{X}_q)^\top \mathbf{K}_n^{-1} \mathbf{k}_n(\mathcal{X}_q) + r(\mathcal{X}_q) \end{aligned}$$

where $\mathbf{k}_n(\mathcal{X}_q) = (k(\mathbf{x}_i, \mathbf{x}'_j))_{1 \leq i \leq n, 1 \leq j \leq q}$ and $\mathbf{K}_n = (k(\mathbf{x}_i, \mathbf{x}_j) + \delta_{i=j} r(\mathbf{x}_i))_{1 \leq i, j \leq n}$. We refer to (Rasmussen and Williams, 2006; Forrester et al., 2008; Ginsbourger, 2018; Gramacy, 2020) and references therein for additional details on GPs and associated sequential design strategies.

Noise variance, $r(\mathbf{x})$, if present, is seldom known and must be estimated. With replication,

stochastic kriging (Ankenman et al., 2010) relies on empirical variance estimates. Otherwise, estimation methods have been proposed, building on the Markov chain Monte Carlo method of (Goldberg et al., 1998), as discussed, e.g., by (Binois et al., 2018). Not only is replication beneficial in terms of variance estimation, it also has an impact on the computational speed of using GPs, where the costs scale with the number of unique designs rather than the total number of evaluations.

2.2 Batch Bayesian optimization

Among many acquisition functions α , we focus on EI, with its analytical expression, compared with, say, entropy criteria. EI (Mockus et al., 1978) is defined as: $\alpha_{EI}(\mathbf{x}) := \mathbb{E}[\max(0, T - Y_n(\mathbf{x}))]$ where T is the best value observed so far in the deterministic case. In the noisy one, taking T as the best mean estimation over sampled designs (Villemonteix et al., 2009a) or the entire space (Gramacy and Lee, 2011), are simple alternatives. Integrating out noise uncertainty is done by Letham et al. (2018), losing analytical tractability.

This acquisition function can be extended to take into account the addition of $q \geq 1$ new points, e.g., with the multi-point EI, $\alpha_{qEI}(\mathcal{X}_q) := \mathbb{E}[\max(0, T - Y_n(\mathbf{x}'_1), \dots, T - Y_n(\mathbf{x}'_q))]$ that has an expression amenable for computation (Chevalier and Ginsbourger, 2013). This is generally not the case for other batch extensions, that can nevertheless be estimated efficiently by Monte Carlo (Wilson et al., 2018). The generic multi-point BO loop is drafted in Algorithm 1.

Algorithm 1 Pseudo-code for batch BO

Require: N_{max} (total budget), q (batch size), GP model trained on initial DoE $(\mathbf{x}_i, y_i)_{1 \leq i \leq n}$

- 1: **while** $n \leq N_{max}$ **do**
 - 2: Choose $\mathbf{x}_{n+1}, \dots, \mathbf{x}_{n+q} \in \arg \max_{\mathcal{X}_q \in \mathbb{X}} \alpha(\mathcal{X}_q)$
 - 3: Update the GP model by conditioning on $\{\mathbf{x}_{n+i}, y_{n+i}\}_{1 \leq i \leq q}$.
 - 4: $n \leftarrow n + q$
 - 5: **end while**
-

2.3 Multi-objective BO

The multi-objective optimization problem (MOOP) is to find $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^d} (f_1(\mathbf{x}), \dots, f_p(\mathbf{x}))$, $p \geq 2$. $p > 4$ is generally considered as the *many-objective* setup, with its own set of challenges for BO, see e.g., (Binois et al., 2020). The solutions of a MOOP are the best compromise solutions between objectives, in the Pareto dominance sense. A solution \mathbf{x} is said to be dominated by another \mathbf{x}' if $\forall i, f_i(\mathbf{x}') \leq f_i(\mathbf{x})$ and $f_i(\mathbf{x}') < f_i(\mathbf{x})$ for at least one i . The Pareto set is the set of solutions that are not dominated by any other design in \mathbb{X} ; the Pareto front is its image in the objective space. In the noisy case, we consider the noisy MOOP commonly defined on expectations over objectives (Hunter et al., 2017).

Measured in the objective space, the hypervolume refers to the volume dominated by a set of points relative to a reference point, see Figure 1. It serves both as a performance metric in MOO, see e.g., Audet et al. (2020) and references therein, or to measure improvement in extending EI (Emmerich et al., 2006). The corresponding expected hypervolume improvement (EHI) can be computed in closed form for two or three objectives (Emmerich et al., 2011; Yang et al., 2017), or by sampling, see e.g., Svenson (2011). A batch version is proposed by Daulton et al. (2020).

3 Batch selection as a portfolio problem

We start by highlighting further the practical difficulties of existing methods to deal with batches of more than a few tens of evaluations. For these reasons, we go back to the roots of BO, the exploration/exploitation trade-off and propose a fast alternative that ignores batch optimization difficulties.

3.1 Challenges with massive batches

Let us summarize the issues of existing works to scale with large batches. For qEI, the exact computation makes use of multivariate normal probabilities, whose computation do not scale well with the batch size. A much faster approximation of the batch EI is described in Binois (2015), relying on nested Gaussian approximations of the maximum of two Gaussian variables from Clark (1961). Otherwise, sampling methods, e.g., with the reparameterization trick (Wilson et al., 2018), are largely used, but may be less precise as the batch size increases. Unsurprisingly, batch criteria evaluation costs increase with the batch size.

This would already be troublesome when optimization iterations must be conducted at a fast pace, but is amplified by the difficulty of actually optimizing the acquisition function itself. While earlier works used branch and bounds (Jones et al., 1998) to guarantee optimality in the single point case, multi-start gradient based optimization or evolutionary algorithms are predominantly used. In the batch setting, the size of the optimization problem becomes $q \times d$, a real challenge for multimodal functions, even with the availability of gradients. Wilson et al. (2018) showed that optimizing the batch members one-by-one is sensible, which still requires to solve q d -dimensional global optimization problems. Relying on discrete search spaces bypasses parts of the problem, even though finding the best batch becomes a combinatorial problem. As a result, only sub-optimal solutions are reachable in practice for batch acquisition function optimization. Furthermore, for massive batches, due to the propensity to avoid clusters of points, some batch members may have a very low influence on the acquisition function value. Rather than optimizing, a perhaps even more computationally intense option is to consider the density under EI. That is, to either find local optima and adapt batch size as in Nguyen et al. (2016), or sampling uniformly from the EI density with slice sampling and clustering as with Groves and Pyzer-Knapp (2018).

Replication adds another dimension to the problem, depending on whether or not replicating is better (or marginally worse) than adding a new design, especially with future

computational speed in mind. Even fixing the number of replicates per batch, the associated dynamic programming optimization problem is hard to solve. Ideas have been developed in the context of global surrogate accuracy, see e.g., Zhang et al. (2020), with the use of back-tracking.

3.2 Exploration/exploitation trade-off

At the core of BO is the idea that regions of interest have either a low mean, or have a large uncertainty given by the variance. Hence EI just provides a way to select a specific solution on the corresponding Pareto front. It was noticed by Jones et al. (1998) in showing that $\frac{\partial EI}{\partial m_n}(x) = -\Phi\left(\frac{T-m_n(\mathbf{x})}{s_n(\mathbf{x})}\right) < 0$ and $\frac{\partial EI}{\partial s_n}(x) = \phi\left(\frac{T-m_n(\mathbf{x})}{s_n(\mathbf{x})}\right) > 0$. Note that for UCB (Srinivas et al., 2010), $\alpha_{UCB}(\mathbf{x}) := m_n(\mathbf{x}) - \sqrt{\beta}\sigma_n(\mathbf{x})$, the tuning parameter β is a way to select one solution on the convex parts of the mean/variance Pareto front.

Similarly for EI, this can be done by taking expectations of powers of the improvement, i.e., the generalized EI (GEI) (Schonlau et al., 1998; Wang et al., 2017), for which higher powers of EI reward larger variance and make it more global, as in Ponweiser et al. (2008). Note that the probability of improvement (PI), the zeroth-order EI, is not on the trade-off Pareto front. This could explain why PI is often discarded as being too exploitative (higher variance is detrimental as soon as the predicted mean is below T).

Our main point is thus that rather than having to define a specific trade-off *a priori*, it is better to find the set of optimal trade-offs and select batch points on it *a posteriori*. When noise is present, we include an additional objective of variance reduction. That is, for two designs with the same mean and variance, the one for which adding an observation will decrease the variance the most is preferred.

3.3 Proposition with HSRI

Yevseyeva et al. (2014) defined the hypervolume Sharpe ratio indicator (HSRI) to select individuals in evolutionary MO algorithms, with further study on their properties in Guerreiro and Fonseca (2016, 2020). In a parallel with portfolio-selection, individual performance is related to expected return, while diversity is related to the return covariance. Here the objectives considered are the mean and standard deviation (rather than the variance, to be on the same scale).

Let $A = \{a^{(1)}, \dots, a^{(r)}\}$ be a non-empty set of assets, let vector $\mathbf{r} \in \mathbb{R}^r$ denote their expected return and matrix $\mathbf{Q} \in \mathbb{R}^{r \times r}$ denote the return covariance between pairs of assets. Let $\mathbf{z} \in [0, 1]^r$ be the investment vector where z_i denotes the investment in asset $a^{(i)}$. The Sharpe-ratio maximization problem is defined as

$$\max_{\mathbf{z} \in [0, 1]^r} h(\mathbf{z}) = \frac{\mathbf{r}^\top \mathbf{z} - r_f}{\sqrt{\mathbf{z}^\top \mathbf{Q} \mathbf{z}}} \quad s.t. \quad \sum_{i=1}^r z_i = 1,$$

with r_f the return of a riskless asset and h the Sharpe ratio. This problem, restated as a convex quadratic programming problem (QP), can be solved efficiently (and only once per iteration). The outcome is a set of weights, from which q evaluations must be selected. In the noiseless case, the q largest weights are selected. When replication is possible, the allocation becomes proportional to the weights: find γ s.t. $\sum_{i=1}^r \lceil \gamma \times z_i^* \rceil = q$, by dichotomy. Ties are resolved randomly.

HSRI is an instance of portfolio selection where the expected return and return covariance are based on the hypervolume improvement, see Figure 1. Importantly, as shown in Guerreiro and Fonseca (2016), if a set of assets is dominated by another set, its Sharpe ratio is lower. Furthermore, no allocation is made on dominated points. Finally they show that only the reference point R needs to be set in practice.

We illustrate the method on an example in Figure 2, comparing the outcome of optimizing

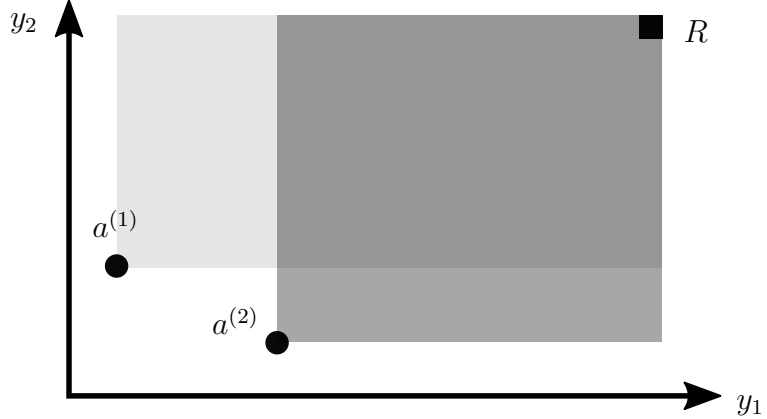


Figure 1: Hypervolume dominated by two assets $a^{(1)}$ (light gray) and $a^{(2)}$ (gray) with respect to the reference point R , corresponding to the expected return. The covariance return is given by the volume jointly dominated by both points (dark gray).

qEI, its fast approximation qAEI (Binois, 2015) and qHSRI in the mean vs. standard deviation space. Additional candidates are shown, either randomly sampled in \mathbb{X} or on the estimated exploration/exploitation Pareto set. Negation of the standard deviation is taken to consider minimization. While all qHSRI selected designs are on the Pareto front, this is not the case for the qEI version, particularly so when q is larger, where none of the selected designs are – possibly due to the much higher difficulty of solving the corresponding optimization problem. Designs corresponding to powers of EI also appear on the Pareto front, showing a richer exploration/exploitation trade-off than with EI only. We also observe that points with large variance may not be of interest if they have a negligible PI.

It could be argued that this approach to batch optimization ignores distances in the input space and could form clusters. While this is the case, since the batch points cover the exploration/exploitation Pareto front, it automatically adapts to this latter’s range of admissible values, depending on the iteration and problem. This is much harder to adapt *a priori* in the input space and it avoids having to define minimal distances manually, as in Gonzalez et al. (2016). Still, for numerical stability and to favor replication, a minimal distance can be added as well.

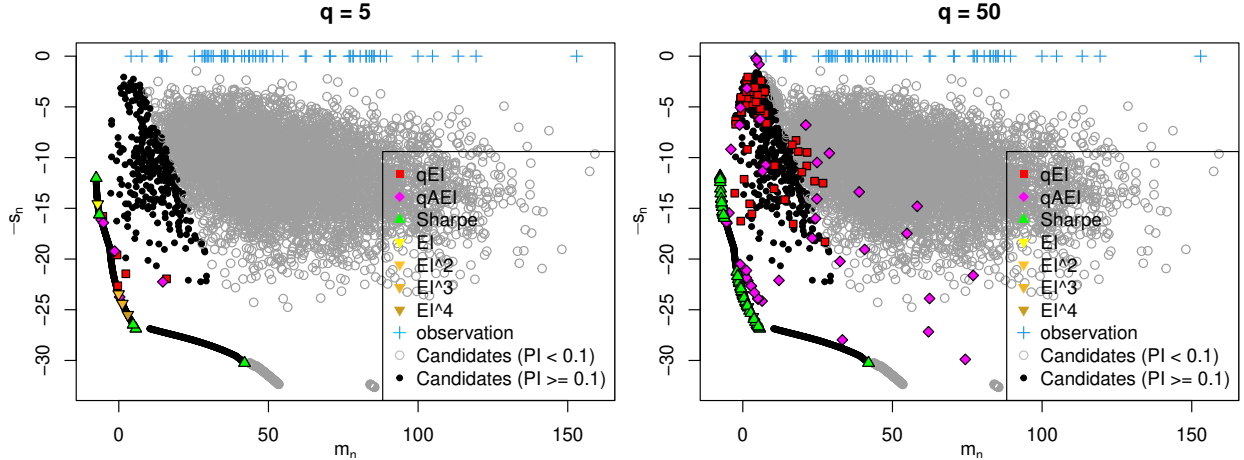


Figure 2: Comparison of qHSRI with qEI and qAEI acquisition functions on a deterministic toy example. The first four GEI optimal solutions are depicted as well.

3.4 Extension to the multi-objective setup

To extend to MO, we propose to search trade-offs on the objective means and averaged predictive standard deviations Pareto front, $\bar{\sigma}_n(\mathbf{x}) = p^{-1} \sum_{i=1}^p s_n^{(i)}(x)/\sigma_n^{(i)}$ with $\sigma_n^{(i)}$ the i^{th} -objective GP variance hyperparameter. Taking all p standard deviations is possible, but in most cases the corresponding objectives are very correlated since they increase with the distance to observed design points. In the case where the GP hyperparameters are the same and evaluations of objectives coupled, the objectives would be perfectly correlated.

The pseudo-code of the approach is given in Algorithm 2. In terms of complexity, the main task is to find the mean/standard deviation Pareto front, where both are $\mathcal{O}(n^2)$ after a $\mathcal{O}(n^3)$ matrix inversion operation that only needs to be done once. An order of magnitude can be gained with approximations, see for instance Wang et al. (2019). Filtering solutions reduces the size of the QP problem to be solved, either with PI or the probability of non-domination (PND). Importantly, the complexity does not depend on q .

Algorithm 2 Pseudo-code for batch BO with qHSRI

Require: N_{max} (total budget), q (batch size), p GP models fitted on initial DoE $(\mathbf{x}_i, y_i)_{1 \leq i \leq n}$

```
1: while  $n \leq N_{max}$  do
2:   Find  $\mathbb{X}_s \in \arg \min m_n^{(1)}(\mathbf{x}), \dots, m_n^{(p)}(\mathbf{x}), \bar{\sigma}_n(\mathbf{x})$ 
3:   Filter dominated solutions in  $\mathbb{X}_s$ 
4:   if  $p = 1$  and  $s > q$  then
5:     Filter points with low PI in  $\mathbb{X}_s$ 
6:   else
7:     Filter points with low PND in  $\mathbb{X}_s$ 
8:   end if
9:   If  $r(x) > 0$ : add variance reduction to objectives
10:  Compute return  $\mathbf{r}$  and covariance matrix  $\mathbf{Q}$ 
11:  Compute optimal Sharpe ratio  $\mathbf{z}^*$ 
12:  Allocate  $q$  points based on the weights
13:  Update the GP model with  $\{\mathbf{x}_{n+i}, y_{n+i}\}_{1 \leq i \leq q}$ .
14:   $n \leftarrow n + q$ 
15: end while
```

4 Experiments

For GP modeling of deterministic tests, we use `DiceKriging` (Roustant et al., 2012) and for noisy ones `hetGP` (Binois and Gramacy, 2021). As we use the same GP models, the comparison shows the effect of the acquisition function choice: qEI or qEHI vs. qHSRI. qEI is either the exact version (Chevalier and Ginsbourger, 2013) in `DiceOptim` (Picheny et al., 2021), or the approximated one from Binois (2015), qAEI. qEI is not run with $q > 10$ due to its much higher computational overhead. qEHI from `GPareto` (Binois and Picheny, 2019) is estimated by Monte Carlo. Random search optimization (RO) provides a baseline.

Optimization of the acquisition functions is performed by a combination of random search, local optimization and evolutionary methods. That is, for qEI, $n_u = 100d$ designs are uniformly sampled in the design space before computing their univariate EI. Then $n_b = 100d$ candidate batches are created by sampling these designs with weights given by EI. Then the best batch for qEI is optimized locally with L-BFGS. A global search with `pso` (Bendtsen, 2012) (population of size 200) is conducted as well and the best batch is

selected. The same principle is applied for qEHI. As for qHSRI, in addition to the same uniform sampling strategy, NSGA-II (Deb et al., 2002) from `mco` (Mersmann, 2020) is used to find the exploration/exploitation compromise surface (population of size 500). The reference point for HSRI computations is obtained from the range over each component, increased by 20%. The R code (R Core Team, 2021) of the approach is packaged and available as supplementary material.

For one objective, the optimality gap, i.e., the difference to a reference best solution, is monitored. For the noisy case, the optimality gap is computed either on noiseless values (when known) or on the estimated minimum over iterations, which is the only element accessible in real applications. The hypervolume metric is used in the MO case, from a reference Pareto front computed using NSGA-II and all designs found by the different methods. We show first that qHSRI works similar to EI and EHI, at a much lower computational cost, before focusing on the motivating application for massive batching.

4.1 Mono-objective results

We consider two standard test problems from the literature, the Branin and Hartmann6 test functions, see e.g., (Roustant et al., 2012). In the deterministic setup, the dimension is increased to accommodate larger batches by taking repeated versions of these problem as used, e.g., in Oh et al. (2018). For the noisy Branin (resp. Hartmann6), we take the first objective of the P1 test function (Parr, 2012) (resp. repeated Hartmann3 (Roustant et al., 2012)) as input standard deviation $r(\mathbf{x})^{\frac{1}{2}}$.

From the results in Figure 3, the performance is similar between qEI, qAEI and qHSRI, except that qHSRI is orders of magnitude faster from Table 1. In the noisy problems, looking at the true optimality gap for observed designs shows good performance of RO, since especially in small dimension like for Branin ($d = 2$), there is a high probability of sampling close to local minima. Nevertheless, the optimality gap of the estimated best solution at each

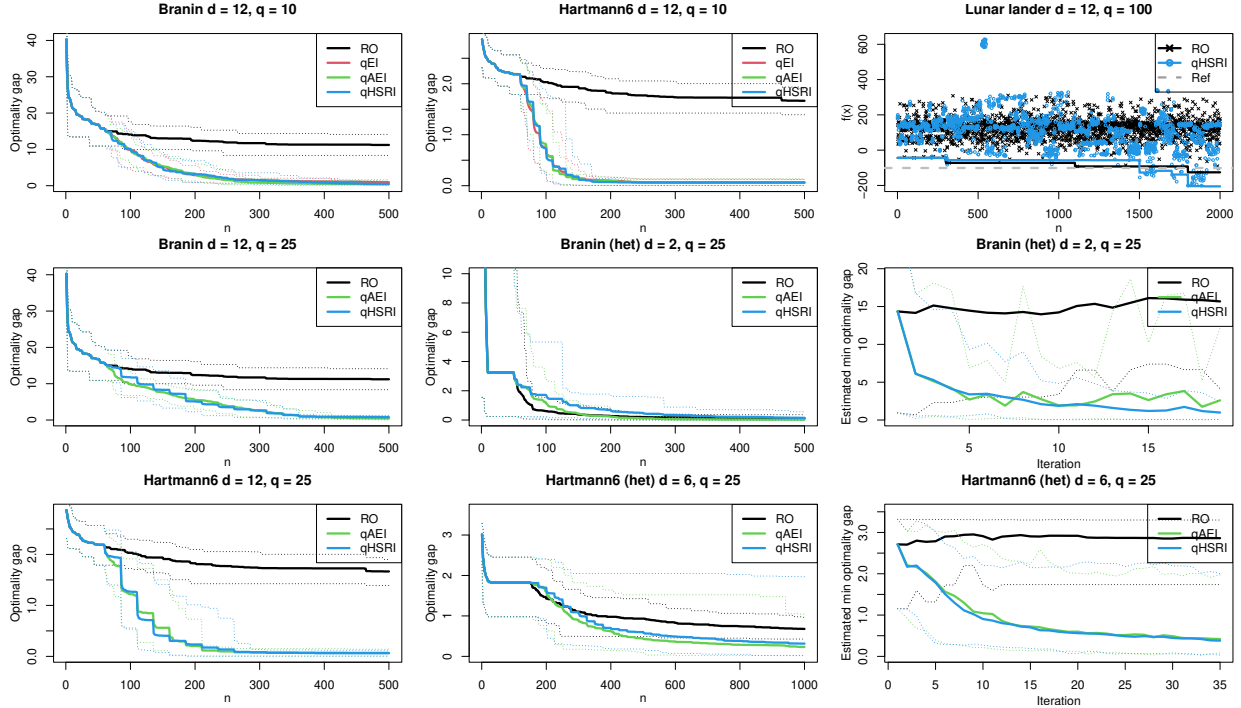


Figure 3: Mono-objective results. Optimality gap for noiseless (resp. noisy) tests over 20 (resp. 40) macro-runs is given, plus optimality gap of estimated minimum for noisy tests. The Lunar lander problem (one single run) shows the evaluated values and estimated minimum found. Thin dashed lines are 5% and 95% quantiles.

iteration is much worse. In addition, the estimated minimum can be slightly better with qHSRI, due to the ability to replicate and thus estimate the noise variance better, which also results in higher speed-ups.

Then we tackle the $12d$ Lunar lander problem (Eriksson et al., 2019). We take $N_{max} = 2000$ with $q = 100$, where a single evaluation is taken as the average over 100 runs to cope with the low signal-to-noise ratio (rather than fixing 50 seeds as in Eriksson et al. (2019)). The solution found (predicted by the GP) is of -205.32 while the reference handcrafted solution gives -101.13 .

4.2 Multi-objective results

We consider the P1 (Parr, 2012) and P2 (Poloni et al., 2000) problems, possibly repeated to increase the dimension. For the noisy setup, one problem serves as the other one’s noise standard deviation function (taking absolute values for positiveness). The hypervolume indicator is used to compare the results, either on the evaluated designs or on the estimated Pareto set (noisy case) at each iteration. The results are shown in Figure 4. Again the performance of qHSRI is on par with qEHI, while much faster (see Table 1) and giving more reliable Pareto front estimation when there is noise.

$f-d-q$	qEI	qAEI	qEHI	qHSRI
B-12-10	24674	5878	–	1120
B-12-25	–	5855	–	458
H-12-10	24234	4941	–	960
H-12-25	–	4575	–	398
B(het)-2-25	–	2792	–	63
H(het)-6-25	–	24349	–	652
P1-6-50	–	–	16042	725
P1(het)-2-50	–	–	8536	191
P2-6-50	–	–	13809	627
P2(het)-2-50	–	–	9016	178

Table 1: Averaged timing results (in s). B is for Branin, H for Hartmann6, and ‘–’ when non applicable.

4.3 CityCOVID data set

We showcase a motivating application example for massive batching: the calibration of the CityCOVID ABM of the population of Chicago in the context of the COVID-19 pandemic, presented in Ozik et al. (2021). The problem is formulated as a nine variable bi-objective optimization problem: the aggregated difference for two target quantities. To inform public health policy, many simulations are necessary in a short period of time, which can only be achieved by running many concurrently. One simulation takes ≈ 10 min, with a low signal-to-

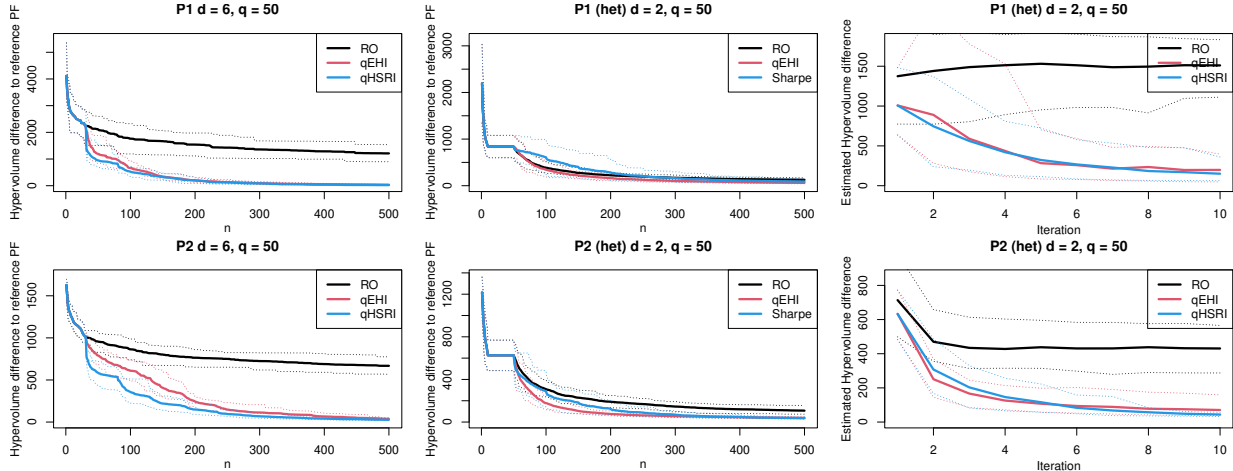


Figure 4: Multi-objective results. The first column is for noiseless tests over 20 macro-runs, the others are for noisy tests over 40 macro-runs, with, in addition to the hypervolume difference over iterations, its estimated counterpart over iterations. Thin dashed lines are 5% and 95% quantiles.

noise ratio. A data set of 217,078 simulations (over 8,368 unique designs, with a degree of replication between 1 and 1000) has been collected by various strategies: IMABC (Rutter et al., 2019), qEHI with fixed degree of replication, and replicating current non-dominated solutions.

From this data set, starting from a subset of the observations (50,585 simulations over 5075 unique designs, with a degree of replication between 3 and 10), qHSRI is used to select candidates among remaining designs up to $q = 2500$ if enough replicates are available, hence with a flexible batch size. To speed up prediction and benefit from a parallel architecture, local GPs are built from 20 nearest neighbors rather than relying on a single global GP. We show in Figure 5 the progress in terms of symmetric difference to the final estimate of the Pareto front, thus penalizing both under and over confident predictions. qHSRI quickly converges to the reference Pareto front, compared to randomly selecting designs.

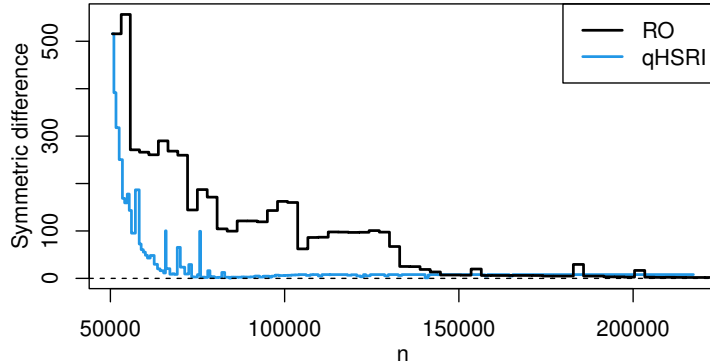


Figure 5: Results on CityCOVID data set.

5 Conclusions and perspectives

Massive batching for BO comes with the additional challenge that batch selection must happen at a fast pace. Here we demonstrate qHSRI as a flexible and light-weight option, independent of the batch size. It also handles replication natively, resulting in additional speed-up for noisy simulators without fixing a degree of replication. Hence this proposed approach makes a sensible, simple, yet efficient, baseline for massive batching.

Possible extensions include taking into account global effects (e.g., on entropy, integrated variance, etc.) of candidate designs to be less myopic. A more dynamic Sharpe ratio allocation could be beneficial as well to improve replication. Finally, while the integration of a few constraints is straightforward, handling more could rely on the use of copulas as in Eriksson and Poloczek (2021) to alleviate the increase of the dimension of the exploration/exploitation trade-off surface. The study of the convergence, e.g., based on results for UCB with various β s, is left for future works.

Acknowledgments

This research was supported by the U.S. Department of Energy (DOE) Office of Science through the National Virtual Biotechnology Laboratory, a consortium of DOE national laboratories focused on response to COVID-19, with funding provided by the Coronavirus

CARES Act. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. The authors are also grateful to the OPAL infrastructure from Université Côte d’Azur for providing resources and support.

References

- Alvi, A., Ru, B., Calliess, J.-P., Roberts, S., and Osborne, M. A. (2019). Asynchronous batch Bayesian optimisation with improved local penalisation. In *International Conference on Machine Learning*, pages 253–262.
- Ankenman, B., Nelson, B. L., and Staum, J. (2010). Stochastic kriging for simulation metamodeling. *Operations research*, 58(2):371–382.
- Audet, C., Bignon, J., Cartier, D., Le Digabel, S., and Salomon, L. (2020). Performance indicators in multiobjective optimization. *European journal of operational research*.
- Azimi, J., Fern, A., and Fern, X. Z. (2010). Batch Bayesian optimization via simulation matching. In *Advances in Neural Information Processing Systems*, pages 109–117. Citeseer.
- Baker, E., Barbillon, P., Fadikar, A., Gramacy, R. B., Herbei, R., Higdon, D., Huang, J., Johnson, L. R., Mondal, A., Pires, B., et al. (2020). Stochastic simulators: An overview with opportunities. *arXiv preprint arXiv:2002.01321*.
- Bendtsen, C. (2012). *pso: Particle Swarm Optimization*. R package version 1.0.3.
- Binois, M. (2015). *Uncertainty quantification on Pareto fronts and high-dimensional strategies in Bayesian optimization, with applications in multi-objective automotive design*. PhD thesis, Mines Saint-Etienne, EMSE.

- Binois, M. and Gramacy, R. B. (2021). hetGP: Heteroskedastic Gaussian process modeling and sequential design in R. *Journal of Statistical Software*, 98(13):1–44.
- Binois, M., Gramacy, R. B., and Ludkovski, M. (2018). Practical heteroscedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 27(4):808–821.
- Binois, M., Huang, J., Gramacy, R. B., and Ludkovski, M. (2019). Replication or exploration? Sequential design for stochastic simulation experiments. *Technometrics*, 61(1):7–23.
- Binois, M. and Picheny, V. (2019). GPareto: An R package for Gaussian-process-based multi-objective optimization and analysis. *Journal of Statistical Software*, 89(8):1–30.
- Binois, M., Picheny, V., Taillandier, P., and Habbal, A. (2020). The Kalai-Smorodinsky solution for many-objective Bayesian optimization. *Journal of Machine Learning Research*, 21(150):1–42.
- 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.
- Chevalier, C. and Ginsbourger, D. (2013). Fast computation of the multi-points expected improvement with applications in batch selection. In *International Conference on Learning and Intelligent Optimization*, pages 59–69. Springer.
- Clark, C. E. (1961). The greatest of a finite set of random variables. *Operations Research*, 9(2):145–162.
- Daulton, S., Balandat, M., and Bakshy, E. (2020). Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems*, 33.

- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evolutionary Computation, IEEE Transactions on*, 6(2):182–197.
- Desautels, T., Krause, A., and Burdick, J. W. (2014). Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *The Journal of Machine Learning Research*, 15(1):3873–3923.
- Emmerich, M., Giannakoglou, K., and Naujoks, B. (2006). Single-and multiobjective evolutionary optimization assisted by Gaussian random field metamodels. *Evolutionary Computation, IEEE Transactions on*, 10(4):421–439.
- Emmerich, M. T. and Deutz, A. H. (2018). A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Natural computing*, 17(3):585–609.
- Emmerich, M. T., Deutz, A. H., and Klinkenberg, J. W. (2011). Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *Evolutionary Computation (CEC), 2011 IEEE Congress on*, pages 2147–2154. IEEE.
- Emmerich, M. T., Yang, K., and Deutz, A. H. (2020). Infill criteria for multiobjective Bayesian optimization. In *High-Performance Simulation-Based Optimization*, pages 3–16. Springer.
- Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M. (2019). Scalable global optimization via local Bayesian optimization. *Advances in Neural Information Processing Systems*, 32:5496–5507.
- Eriksson, D. and Poloczek, M. (2021). Scalable constrained Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 730–738. PMLR.

- Forrester, A., Sobester, A., and Keane, A. (2008). *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.
- Frazier, P. I. (2018). Bayesian optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*, pages 255–278. INFORMS.
- Ginsbourger, D. (2018). *Sequential Design of Computer Experiments*, pages 1–9. American Cancer Society.
- Ginsbourger, D., Le Riche, R., and Carraro, L. (2010). Kriging is well-suited to parallelize optimization. In *Computational Intelligence in Expensive Optimization Problems*, pages 131–162. Springer.
- Goldberg, P. W., Williams, C. K., and Bishop, C. M. (1998). Regression with input-dependent noise: A Gaussian process treatment. *Advances in neural information processing systems*, pages 493–499.
- Gonzalez, J., Dai, Z., Hennig, P., and Lawrence, N. (2016). Batch Bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 648–657.
- Gonzalez, S. R., Jalali, H., and Van Nieuwenhuysse, I. (2020). A multiobjective stochastic simulation optimization algorithm. *European Journal of Operational Research*, 284(1):212–226.
- Gramacy, R. B. (2020). *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. CRC Press.
- Gramacy, R. B. and Lee, H. K. (2009). Adaptive design and analysis of supercomputer experiments. *Technometrics*, 51(2):130–145.

- Gramacy, R. B. and Lee, H. K. H. (2011). Optimization under unknown constraints. In Bernardo, J., Bayarri, S., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 9*, pages 229–256. Oxford University Press.
- Groves, M. and Pyzer-Knapp, E. O. (2018). Efficient and scalable batch Bayesian optimization using K-means. *arXiv preprint arXiv:1806.01159*.
- Guerreiro, A. P. and Fonseca, C. M. (2016). Hypervolume Sharpe-ratio indicator: Formalization and first theoretical results. In *International Conference on Parallel Problem Solving from Nature*, pages 814–823. Springer.
- Guerreiro, A. P. and Fonseca, C. M. (2020). An analysis of the Hypervolume Sharpe-Ratio Indicator. *European Journal of Operational Research*, 283(2):614–629.
- Haftka, R. T., Villanueva, D., and Chaudhuri, A. (2016). Parallel surrogate-assisted global optimization with expensive functions—a survey. *Structural and Multidisciplinary Optimization*, 54(1):3–13.
- Hennig, P. and Schuler, C. J. (2012). Entropy search for information-efficient global optimization. *The Journal of Machine Learning Research*, 98888:1809–1837.
- Hernández-Lobato, J. M., Requeima, J., Pyzer-Knapp, E. O., and Aspuru-Guzik, A. (2017). Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In *International conference on machine learning*, pages 1470–1479. PMLR.
- Hoffman, M. D., Brochu, E., and de Freitas, N. (2011). Portfolio allocation for Bayesian optimization. In *UAI*, pages 327–336. Citeseer.
- Hunter, S. R., Applegate, E. A., Arora, V., Chong, B., Cooper, K., Rincón-Guevara, O., and Vivas-Valencia, C. (2017). An introduction to multi-objective simulation optimization. *Optimization Online*.

- Janusevskis, J., Le Riche, R., Ginsbourger, D., and Girdziusas, R. (2012). Expected improvements for the asynchronous parallel global optimization of expensive functions: Potentials and challenges. In *Learning and Intelligent Optimization*, pages 413–418. Springer.
- Jones, D., Schonlau, M., and Welch, W. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492.
- Kandasamy, K., Krishnamurthy, A., Schneider, J., and Póczos, B. (2018). Parallelised Bayesian optimisation via Thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142. PMLR.
- Letham, B., Karrer, B., Ottoni, G., Bakshy, E., et al. (2018). Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis*.
- Lukovic, M. K., Tian, Y., and Matusik, W. (2020). Diversity-guided multi-objective Bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems*, 33:6–12.
- Marmin, S., Chevalier, C., and Ginsbourger, D. (2015). Differentiating the multipoint expected improvement for optimal batch design. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 37–48. Springer.
- Mersmann, O. (2020). *mco: Multiple Criteria Optimization Algorithms and Related Functions*. R package version 1.15.6.
- Mockus, J., Tiesis, V., and Zilinskas, A. (1978). The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2(117-129):2.
- Nguyen, V., Rana, S., Gupta, S. K., Li, C., and Venkatesh, S. (2016). Budgeted batch Bayesian optimization. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 1107–1112. IEEE.

- Oh, C., Gavves, E., and Welling, M. (2018). Bock: Bayesian optimization with cylindrical kernels. In *International Conference on Machine Learning*, pages 3868–3877. PMLR.
- Ozik, J., Wozniak, J. M., Collier, N., Macal, C. M., and Binois, M. (2021). A population data-driven workflow for COVID-19 modeling and learning. *The International Journal of High Performance Computing Applications*, 35(5):483–499.
- Parr, J. M. (2012). *Improvement Criteria for Constraint Handling and Multiobjective Optimization*. PhD thesis, University of Southampton.
- Picheny, V., Green, D. G., and Roustant, O. (2021). *DiceOptim: Kriging-Based Optimization for Computer Experiments*. R package version 2.1.1.
- Poloni, C., Giurgevich, A., Onesti, L., and Pediroda, V. (2000). Hybridization of a multi-objective genetic algorithm, a neural network and a classical optimizer for a complex design problem in fluid dynamics. *Computer Methods in Applied Mechanics and Engineering*, 186(2):403–420.
- Ponweiser, W., Wagner, T., and Vincze, M. (2008). Clustered multiple generalized expected improvement: A novel infill sampling criterion for surrogate models. In *Proc. (IEEE World Congress Computational Intelligence). IEEE Congress Evolutionary Computation CEC 2008*, pages 3515–3522.
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rasmussen, C. E. and Williams, C. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Rojas-Gonzalez, S. and Van Nieuwenhuysse, I. (2020). A survey on kriging-based infill

- algorithms for multiobjective simulation optimization. *Computers & Operations Research*, 116:104869.
- Rontsis, N., Osborne, M. A., and Goulart, P. J. (2020). Distributionally robust optimization techniques in batch Bayesian optimization. *Journal of Machine Learning Research*, 21(149):1–26.
- Roustant, O., Ginsbourger, D., and Deville, Y. (2012). DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization. *Journal of Statistical Software*, 51(1):1–55.
- Rutter, C. M., Ozik, J., DeYoreo, M., and Collier, N. (2019). Microsimulation model calibration using incremental mixture approximate Bayesian computation. *The Annals of Applied Statistics*, 13(4):2189–2212.
- Schonlau, M., Welch, W. J., and Jones, D. R. (1998). Global versus local search in constrained optimization of computer models. *Lecture Notes-Monograph Series*, pages 11–25.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.
- Sóbestor, A., Leary, S. J., and Keane, A. J. (2004). A parallel updating scheme for approximating and optimizing high fidelity computer simulations. *Structural and multidisciplinary optimization*, 27(5):371–383.
- 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 International Conference on Machine Learning*, pages 1015–1022.

- Svenson, J. D. (2011). *Computer Experiments: Multiobjective Optimization and Sensitivity Analysis*. PhD thesis, The Ohio State University.
- Tran, A., Sun, J., Furlan, J. M., Pagalthivarthi, K. V., Visintainer, R. J., and Wang, Y. (2019). pBO-2GP-3B: A batch parallel known/unknown constrained Bayesian optimization with feasibility classification and its applications in computational fluid dynamics. *Computer Methods in Applied Mechanics and Engineering*, 347:827–852.
- Villemonteix, J., Vazquez, E., Sidorkiewicz, M., and Walter, E. (2009a). Global optimization of expensive-to-evaluate functions: an empirical comparison of two sampling criteria. *Journal of Global Optimization*, 43(2):373–389.
- Villemonteix, J., Vazquez, E., and Walter, E. (2009b). An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509–534.
- Wang, H., van Stein, B., Emmerich, M., and Back, T. (2017). A new acquisition function for Bayesian optimization based on the moment-generating function. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 507–512. IEEE.
- Wang, J., Clark, S. C., Liu, E., and Frazier, P. I. (2020). Parallel Bayesian global optimization of expensive functions. *Operations Research*, 68(6):1850–1865.
- Wang, K., Pleiss, G., Gardner, J., Tyree, S., Weinberger, K. Q., and Wilson, A. G. (2019). Exact Gaussian processes on a million data points. *Advances in Neural Information Processing Systems*, 32:14648–14659.
- Wang, Z., Gehring, C., Kohli, P., and Jegelka, S. (2018). Batched large-scale Bayesian optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence and Statistics*.

- Wang, Z. and Jegelka, S. (2017). Max-value entropy search for efficient Bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR.
- Wilson, J., Hutter, F., and Deisenroth, M. (2018). Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 9884–9895.
- Yang, K., Emmerich, M., Deutz, A., and Fonseca, C. M. (2017). Computing 3-D expected hypervolume improvement and related integrals in asymptotically optimal time. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 685–700. Springer.
- Yevseyeva, I., Guerreiro, A. P., Emmerich, M. T., and Fonseca, C. M. (2014). A portfolio optimization approach to selection in multiobjective evolutionary algorithms. In *International Conference on Parallel Problem Solving from Nature*, pages 672–681. Springer.
- Zhang, B., Gramacy, R. B., Johnson, L., Rose, K. A., and Smith, E. (2020). Batch-sequential design and heteroskedastic surrogate modeling for delta smelt conservation. *arXiv preprint arXiv:2010.06515*.
- Zhang, J., Ma, Y., Yang, T., and Liu, L. (2017). Estimation of the Pareto front in stochastic simulation through stochastic kriging. *Simulation Modelling Practice and Theory*, 79:69–86.
- Zhang, Q., Liu, W., Tsang, E., and Virginas, B. (2010). Expensive multiobjective optimization by MOEA/D with Gaussian process model. *Evolutionary Computation, IEEE Transactions on*, 14(3):456–474.
- Zuluaga, M., Krause, A., and Püschel, M. (2016). e-pal: An active learning approach to the multi-objective optimization problem. *Journal of Machine Learning Research*, 17:1–32.