

# A portfolio approach to massively parallel Bayesian optimization

Anonymous authors

Paper under double-blind review

## 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 to simultaneously select multiple designs via an infill criterion. Still, despite the increased availability of computing resources that enable large-scale parallelism, the strategies that work for selecting a few tens of parallel designs for evaluations become limiting due to the complexity of selecting more designs. 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 noisy functions, for mono and multi-objective optimization tasks. These experiments show orders of magnitude speed improvements over existing methods with similar or better performance.

## 1 Introduction

Current trends in improving the speed or accuracy of individual computations are on exploiting parallelization on highly concurrent computing systems. These computer models (a.k.a. simulators) are prevalent in many fields, ranging from physics to biology and engineering. Still, increasing the parallelization for individual simulators often comes with diminishing returns and model evaluation time remains limiting. A strategy is then to conduct several evaluations simultaneously, in batches, to optimize (here minimize) quantities of interest. See, e.g., (Haftka et al., 2016) for a review.

For fast simulators, evolutionary algorithms (EAs) are amenable to parallelization by design, see e.g., the review by Emmerich and Deutz (2018). 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; Garnett, 2022), is preferred, with its ability to carefully select the next evaluations. Typically, BO relies on a Gaussian process (GP) model of the simulator, or any black-box, by using a probabilistic surrogate model 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). Alternatives 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., 2022), 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 by limiting the number of unique designs, 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. For EI, rather than just looking at its local optima (Sóbester et al., 2004), some heuristics propose to select batch points iteratively, replacing unknown values at selected

points by pseudo-values (Ginsbourger et al., 2010). This was coined as “hallucination” in the UCB version of (Desautels et al., 2014). Exact expressions have been derived for the batch EI (qEI) by (Chevalier and Ginsbourger, 2013) and for its gradient 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). More generally, 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 requires solving a semi-definite problem, limiting the scaling up to large batches. Gonzalez et al. (2016) reduce batch selection cost by not modeling the joint probability distribution of the batch nor using a hallucination scheme. Their idea is to select batch members sequentially by penalizing proximity to the previously selected ones. Taking different infill 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 and not intended as a mechanism to select batch candidates. Using local models is another way 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. For entropy or stepwise uncertainty reduction criteria, see, e.g., (Chevalier et al., 2014a), 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 the calibration of a large-scale agent-based model (ABM) of COVID-19 run on a supercomputer (Ozik et al., 2021). 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, with the added goal of reducing the *time to solution* for meaningful, i.e., timely, decision-making support. See, e.g., (Hunter et al., 2017) for a review of options for black-boxes and (Emmerich et al., 2020) for multi-objective (MO) BO infill criteria. MO versions of batch algorithms have also been proposed, taking different scalarization weights (Zhang et al., 2010), or relying on an additional notion of diversity (Lukovic et al., 2020).

The operations research community has seen works dealing with low signal-to-noise ratios and heteroscedasticity, 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 Nieuwenhuyse, 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. Parallel infill criteria are more expensive to evaluate, and even computational costs increasing linearly in the batch size ( $q$ ) become impractical for hundreds or thousands of batch points. Second, optimizing these parallel infill criteria is increasingly harder with larger  $q$ . That is, frontally optimizing the acquisition function requires the solving of a global optimization problem in dimension number of variables  $d \times q$ . 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. 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 simulation 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 (e.g., HPC cluster or supercomputer) with the ability to run hundreds of simulation evaluations in parallel over several iterations for the purpose of reducing the overall time to solution of the optimization to support rapid turnaround of analyses for high consequence decision making (e.g., public health (Ozik et al., 2021), meteorological (Goubier et al., 2020), and other emergency response (Mandel et al., 2019)). 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 the available computing concurrency). The method we propose is to directly identify candidates realizing different exploration/exploitation trade-offs. This amounts to approximating the GP predictive mean vs. variance Pareto front, which is orders of magnitude faster than optimizing most existing batch infill criteria. In doing so, we shift the paradigm of optimizing (or sampling) acquisition functions over candidate batches to quickly finding a set of desirable candidates

to choose from. In the noisy case, possibly with input-dependent variance, the variance reduction makes an additional objective to further 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 extension is to take the mean and variance of each objective, and still select batch-candidates based on the HSRI. The contributions of this work are: **1)** The use of a portfolio allocation strategy for batch-BO, defined on the exploration/exploitation trade-off. It extends directly to the multi-objective setup; **2)** An approach independent of the size of the batch, removing limitations of current batch criteria for large  $q$ ; **3)** The potential for flexible batch sizes and asynchronous allocation via the portfolio approach; **4)** The ability to natively take into account replication and to cope with input-dependent noise variance.

In Section 2 we briefly present GPs, batch BO and MO BO. In Section 3, the difficulties of batch BO are described before detailing the proposed method. The method 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 of the black-box function  $f$ : find  $\mathbf{x}^* \in \underset{\mathbf{x} \in \mathbb{X} \subseteq \mathbb{R}^d}{\operatorname{argmin}} f(\mathbf{x})$  where  $\mathbb{X}$  is typically a hypercube. Among various options for surrogate modeling of  $f$ , 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  $y_i = f(\mathbf{x}_i) + \varepsilon_i$ ,  $\varepsilon_i \sim \mathcal{N}(0, \tau(\mathbf{x}_i))$ , often obtained with a Latin hypercube sample 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  $\sigma^2$ , 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,  $Y_n(\mathcal{X}_q) | \mathbf{y} \sim \mathcal{N}(m_n(\mathcal{X}_q), s_n^2(\mathcal{X}_q))$ :

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

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  $\mathbf{K}_n = (k(\mathbf{x}_i, \mathbf{x}_j) + \delta_{i=j}\tau(\mathbf{x}_i))_{1 \leq i, j \leq n}$ . 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,  $\tau(\mathbf{x})$ , 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 acquisition functions  $\alpha$ , 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 case, taking  $T$  as the best mean estimation over sampled designs (Villemonteix et al., 2009a) or the entire space (Gramacy and Lee, 2011), are 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  $q$  new points, e.g., with the batch (q in short) 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). 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 by Monte Carlo, e.g., with the reparameterization trick (Wilson et al., 2018), are largely used, but may be less precise as the batch size increases. The generic batch BO loop is illustrated in Appendix A.

### 2.3 Multi-objective BO

The multi-objective optimization problem (MOOP) is to find  $\mathbf{x}^* \in \underset{\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^d}{\operatorname{argmin}}(f_1(\mathbf{x}), \dots, f_p(\mathbf{x}))$ ,  $p \geq 2$ .  $p > 4$  is often called 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 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 of EHI, qEHI, is proposed by Daulton et al. (2020; 2021).

## 3 Batch selection as a portfolio problem

We first summarize issues of BO to handle large batches. We then return to the roots of BO, the exploration/exploitation trade-off, to propose an alternative.

### 3.1 Challenges with massive batches

The exact qEI expression uses multivariate normal probabilities, whose computation do not scale well with the batch size. There are also many approximated criteria for batch EI, or similar criteria. However, in all current approaches, the evaluation costs increase with the batch size, at best linearly in  $q$  for existing criteria, which remains too costly for the regime we target.

This is already troublesome when optimization iterations must be conducted quickly, but is amplified by the difficulty of optimizing the acquisition function itself. While earlier works used branch and bounds (Jones et al., 1998) to guarantee optimality with  $q = 1$ , multi-start gradient based optimization or EAs 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 greedily optimizing 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 issue, even though finding the best batch becomes a combinatorial search. In between the greedy and joint options is the work by Daxberger and Low (2017), to optimize an approximated batch-UCB criterion as a distributed constraint problem. As a result, only sub-optimal solutions are reachable in practice for batch acquisition function optimization. Rather than optimizing, a perhaps even more computationally intensive 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). Thompson sampling for mono-objective batch BO, see e.g., Kandasamy et al. (2018), also bypasses the issues of optimizing the acquisition function, but batch members are independently obtained, which can be wasteful for large batches.

Replication adds another layer of decision: whether or not adding a new design is worth the future extra computational cost, compared to the perhaps marginally worse option of replicating on the acquisition function value. With high noise, choosing the amount of replication becomes important, as individual evaluations contain almost no information. But even fixing the number of replicates per batch, selecting batch design locations plus replication degree makes a hard dynamic programming optimization problem.

### 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 variance. This is the BO exploration/exploitation trade-off, see e.g., Garnett (2022). From a multi-objective point of view, acquisition functions resolve this trade-off by selecting a solution on the corresponding mean vs. standard deviation ( $m_n/s_n$ ) Pareto front  $\mathcal{P}$ . With UCB (Srinivas et al., 2010),  $\alpha_{UCB}(\mathbf{x}) := m_n(\mathbf{x}) - \sqrt{\beta} s_n(\mathbf{x})$ , the tuning parameter  $\beta$  is a way to select one solution on the convex parts of this Pareto front. EI can be interpreted this way as well, as noticed by Jones et al. (1998); De Ath et al. (2021) 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  $\frac{\partial EI}{\partial s_n}(x) = \phi\left(\frac{T-m_n(\mathbf{x})}{s_n(\mathbf{x})}\right) > 0$ , where  $\phi$  (resp.  $\Phi$ ) are the Gaussian pdf (resp. cdf). Hence EI also selects a specific solution on the corresponding Pareto front. Navigating the Pareto front 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), which corresponds to the zeroth-order EI, is not on the trade-off Pareto front  $\mathcal{P}$ , explaining 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 that rather than having to define a specific trade-off between exploration and exploitation *a priori*, before considering batching, it is better to find the set of optimal trade-offs  $\mathcal{P}$  and select batch points from it *a posteriori*.

### 3.3 Proposition with HSRI

We propose to use a criterion to select solutions on the exploration-exploitation Pareto front, rather than doing so randomly as in Gupta et al. (2018). Yevseyeva et al. (2014) defined the hypervolume Sharpe ratio indicator (HSRI) to select individuals in MO EAs, with further study on their properties in Guerreiro and Fonseca (2016; 2020). Here, we borrow from the portfolio-selection approach, where individual performance is related to expected return, while diversity is related to the return covariance.

Let  $A = \{\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(l)}\}$  be a non-empty set of assets,  $\mathbf{a}^{(i)} \in \mathbb{R}^s$ ,  $s \geq 2$ , let vector  $\mathbf{r} \in \mathbb{R}^l$  denote their expected return and matrix  $\mathbf{Q} \in \mathbb{R}^{l \times l}$  denote the return covariance between pairs of assets. Let  $\mathbf{z} \in [0, 1]^l$  be the investment vector where  $z_i$  denotes the investment in asset  $\mathbf{a}^{(i)}$ . The Sharpe-ratio maximization problem is defined as

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

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 only once per iteration. The outcome is a set of weights, from which  $q$  evaluations must be selected. In the noiseless case, they are the  $q$  largest weights. When replication is possible, the allocation becomes proportional to the weights: find  $\gamma$  s.t.  $\sum_{i=1}^l [\gamma \times z_i^*] = q$ , by dichotomy (randomly resolving ties).

HSRI Guerreiro and Fonseca (2016) is an instance of portfolio selection where the expected return and return covariance are based on the hypervolume improvement:  $r_i = p_{ii}$  and  $Q_{ij} = p_{ij} - p_{ii}p_{jj}$  where  $p_{ij} = \left( \prod_{1 \leq t \leq p} (R_t - \max(a_t^{(i)}, a_t^{(j)})) \right) / \left( \prod_{1 \leq t \leq p} (R_t - f_t^*) \right)$ ; see Figure 1 for an illustration. Here, instead of using actual objective values in MO EAs, an asset  $\mathbf{a}^{(i)}$ , corresponding to candidate design  $\mathbf{x}^{(i)}$ , is characterized by its GP predictive mean(s) and standard deviation(s), i.e., with  $p = 1$ ,  $a_1^{(i)} = m_n(\mathbf{x}_i)$  and  $a_2^{(i)} = -s_n(\mathbf{x}^{(i)})$ . Also  $r_f = 0$  since risk-less assets (noiseless observations) bring no improvement. Note that this hypervolume computation scales linearly with the number of objectives. Importantly, as shown in Guerreiro and Fonseca




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).

(2016), if a set of assets is dominated by another set, its Sharpe ratio is lower. Furthermore, no allocation is made on dominated points: they are all on  $\mathcal{P}$ . Finally they show that only the reference point  $R$  needs to be set in practice.

We illustrate the proposed method (qHSRI) on an example in Figure 2, comparing the outcome of optimizing qEI, its fast approximation qAEI (Binois, 2015) and qHSRI in the mean vs. standard deviation ( $m_n$  vs.  $s_n$ ) 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 for minimization. While all qHSRI selected designs are on  $\mathcal{P}$ , 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  $\mathcal{P}$ , 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 (e.g.,  $< 0.1$ ).




Figure 2: Comparison of qHSRI with qEI and qAEI acquisition functions on the noiseless repeated Branin function ( $d = 6$ ,  $n = 60$ ). The first four GEI optimal solutions are depicted as well.

It could be argued that the qHSRI approach ignores distances in the input space and could form clusters. While this is the case, since the batch points cover  $\mathcal{P}$ , it automatically adapts to this latter’s range of optimal values, depending on the iteration and problem. This is 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 set as well.

### 3.4 Extension to the multi-objective setup

To extend to MO, we propose to search trade-offs on the objective means,  $m_n^{(1)}(x), \dots, m_n^{(p)}(x)$ , 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  $(\sigma_n^{(i)})^2$  the  $i^{\text{th}}$ -objective GP variance hyperparameter, a  $p+1$  dimensional space. Taking all  $p$  standard deviations is possible, but the corresponding objectives are 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 1. In terms of complexity, the main task is to find the assets, i.e., candidate designs on  $\mathcal{P}$ . Evaluating  $m_n$  and  $s_n$  cost  $\mathcal{O}(n^2)$  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). Then  $\mathbf{r}$  and  $\mathbf{Q}$  are computed on the assets, before maximizing the Sharpe ratio, whose optimal weights provide the best  $q$  designs. Filtering solutions reduces the size of the QP problem to be solved, either with PI or the probability of non-domination (PND) in the MO case. Crucially, the complexity does not depend on  $q$ .

---

**Algorithm 1** Pseudo-code for batch BO with qHSRI

---

**Input:**  $N_{\max}$  (total budget),  $q$  (batch size),  $p$  GP model(s) 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**  $p = 1$  **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**
- 9:   If  $\tau(\mathbf{x}) > 0$ : 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(s) with  $\{\mathbf{x}_{n+i}, y_{n+i}\}_{1 \leq i \leq q}$ .
- 14:    $n \leftarrow n + q$
- 15: **end while**

---

### 3.5 Replication

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 predictive variance the most is preferred. This decrease is given by GP update equations, see e.g., Chevalier et al. (2014b), and do not depend on the value at the future  $q$  designs:  $s_{n+q}^2(\mathcal{X}_q) = s_n^2(\mathcal{X}_q) - s_n^2(\mathcal{X}_q, \mathbf{x}_{1:(n+q)})(s_n^2(\mathbf{x}_{1:(n+q)}))^{-1}s_n^2(\mathbf{x}_{1:(n+q)}, \mathcal{X}_q)$  with  $\mathbf{x}_{1:(n+q)}$  the current DoE augmented by future  $q$  designs. It does depend on the noise variance and the degree of replication, see, e.g., Binois et al. (2019), which may be used to define a minimal degree of replication at candidate designs to ensure a sufficient decrease. Similarly, limiting the replication degree when the decrease of variance of further replication is too low is doable.

## 4 Experiments

Except Wang et al. (2018) that uses disconnected local GP models and Gupta et al. (2018) that also uses the exploration-exploitation PF, existing batch BO methods mostly give results with low  $q$ , e.g.,  $q \leq 20$ . On the implementations we could test, these criteria take more than seconds per evaluation with  $q \approx 100$ , while, in our approach, predicting for a given design takes less than a millisecond. Consequently, comparisons with qHSRI are not feasible for massive  $q$ .

The R package `hetGP` (Binois and Gramacy, 2021) is used for noisy GP modeling. As we use the same GP model, 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 available for  $q > 20$  nor in the noisy case. In the mono-objective case, we also include Thompson sampling, qTS, implemented by generating GP realisations on discrete sets of size  $200d$ . qEHI is from `GPareto` (Binois and Picheny, 2019), estimated by Monte Carlo. PF denotes the method from Gupta et al. (2018), where the batch members are randomly selected on the exploration-exploitation Pareto front. Random search optimization (RO) is added as a baseline. All start with the same space filling designs of size  $5d$ , replicated five times each to help the heteroscedastic GP modeling with low signal to noise ratios.

Optimization of the acquisition functions is performed by combining random search, local optimization and EAs. 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 corresponding best batch for qEI is optimized locally with L-BFGS-B. A global search with `pso` (Bendtsen, 2012) (population of size 200) is conducted too, to directly optimize qEI, and the overall best qEI batch is selected. The same principle is applied for qEHI. As for qHSRI and PF, in addition to the same uniform sampling strategy with  $n_u$  designs, NSGA-II (Deb et al., 2002) from `mco` (Mersmann, 2020) is used to find  $\mathcal{P}$ , the exploration/exploitation compromise surface (with a population of size 500). The reference point  $R$  for HSRI computations is obtained from the range over each component, extended by 20%, as is the default in `GPareto` Binois and Picheny (2019). The R code (R Core Team, 2023) of the approach is available as supplementary material.

For one objective, the optimality gap, i.e., the difference to a reference solution, is monitored. With noise, the optimality gap is computed both on noiseless values (when known) and 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.

In a first step, for validating qHSRI, we compared it with alternatives for relatively low  $q$  values. These preliminary experiments on noiseless functions are provided in Appendix B. The outcome is that qHSRI gives results on par with qEI and qEHI looking at the performance over iterations, at a much lower computational cost. These results also motivate the use of qAEI as a proxy for qEI when this latter is not available. We notice that qTS performed poorly on these experiments, possibly because using discretized realisations is insufficient for the relatively large input dimension ( $d = 12$ ). There the use of random or Fourier features may help Mutny and Krause (2018). As for PF, it requires more samples than qHSRI, even if it goes as fast. This indicates that the portfolio allocation strategy is beneficial compared to randomly sampling on the exploration-exploitation Pareto front.

#### 4.1 Mono-objective results

We first consider the standard Branin and Hartmann6 test functions, see e.g., (Roustant et al., 2012). 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  $\tau(\mathbf{x})^{\frac{1}{2}}$ , hence with heteroscedastic noise (denoted by *het* below).

Figure 3 highlights that qHSRI is orders of magnitude faster than qTS and qAEI for decreasing the optimality gap (see also Table 1 for all timing results). It also improves over random selection on  $\mathcal{P}$  as with PF. In these 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 one of its three global minima. Also, replicating incurs less exploration, penalizing qHSRI on this metric. Nevertheless, the actual metric of interest is the optimality gap of the estimated best solution at each iteration. It is improved with qHSRI, especially for Branin, while performances of the various acquisition functions are similar iteration-wise in the Hartmann6 case. Concerning speed, part of the speed-ups are due to the ability to replicate. Indeed, as apparent in supplementary Figure 9, for qHSRI the number of unique designs remains low, less than 20% of the total number of observations without degrading the sample efficiency. Notice that






Figure 3: Mono-objective results over iterations and over time. Optimality gap and estimated optimality gap for noisy tests over 40 macro-runs are given. Thin dashed lines are 5% and 95% quantiles.

taking larger batches can be faster since the batch selection is independent of  $q$  with qHSRI. Also there are fewer iterations for the same simulation budget, hence less time is spent in fitting GPs and finding  $\mathcal{P}$ .

Next we tackle the more realistic 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 to make it deterministic as in Eriksson et al. (2019)). The solution found (predicted by the GP) is of  $-205.32$  while the reference handcrafted solution gives  $-101.13$ , see Figure 5. The Lunar lander problem with qHSRI took 5 hours; it did not complete with qAEI even in 24 hours due to  $q = 100$ .

## 4.2 Multi-objective results

We consider the P1 (Parr, 2012) and P2 (Poloni et al., 2000) problems. For the noisy setup, one problem serves as the other one’s noise standard deviation function (taking absolute values for positiveness). The results are shown in Figure 4, where the leftmost panels show the beneficial performance of the qHSRI approach in terms of time to solution. While RO performs relatively well looking solely at the hypervolume difference for the evaluated designs (rightmost panels), this does not relate to the quality of the Pareto front estimation. Indeed, the estimated Pareto with RO, that is, using only the non-dominated noisy observations, is far from the actual Pareto front, due to noise realizations. There the Pareto fronts estimated by GP models do show a convergence to the reference Pareto front, indicating that their estimation improves over iterations (middle panels). Finally, like in the mono-objective results, we demonstrate in Appendix B that for the noiseless case the sample efficiency of qHSRI is at least on par to that of qEHI, and even slightly better in some cases.

## 4.3 CityCOVID data set

We showcase a motivating application example for massive batching: calibrating the CityCOVID ABM of the population of Chicago in the context of the COVID-19 pandemic, presented in Ozik et al. (2021) and built on the ChiSIM framework (Macal et al., 2018). It models the 2.7 million residents of Chicago as they move between 1.2 million places based on their hourly activity schedules. The synthetic population of agents extends an existing general-purpose synthetic population Cajka et al. (2010) and statistically matches Chicago’s demographic composition. Agents colocate in geolocated places, which include households, schools, workplaces, etc. The agent hourly activity schedules are derived from the American Time Use Survey and the Panel Study of Income Dynamics and assigned based on agent demographic characteristics.

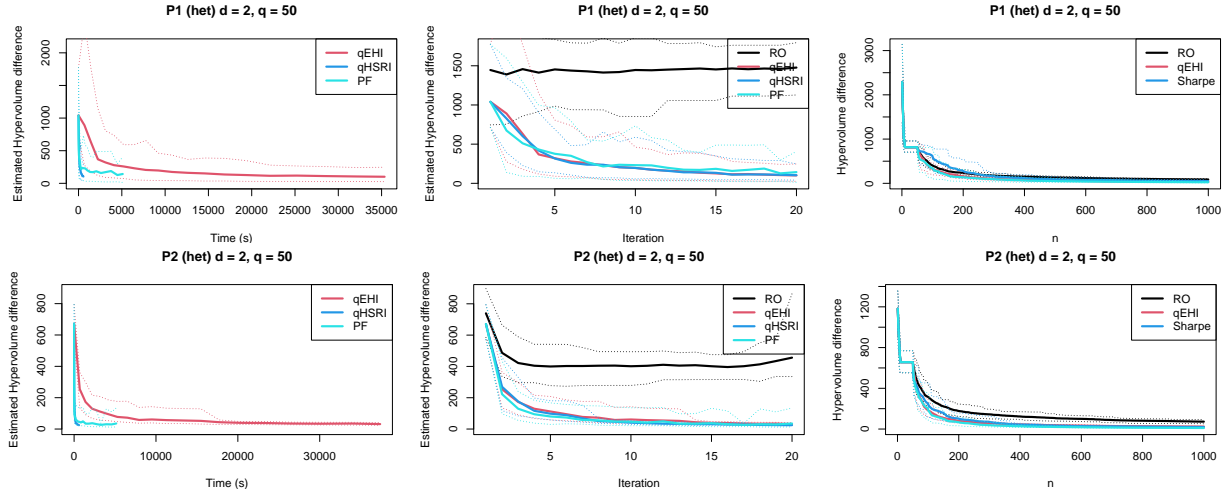


Figure 4: Multi-objective results over time and iterations. Hypervolume difference over a reference Pareto front and its counterpart for the estimated Pareto set for noisy tests over 40 macro-runs are given. Thin dashed lines are 5% and 95% quantiles.

CityCOVID includes COVID-19 disease progression within each agent, including differing symptom severities, hospitalizations, and age-dependent probabilities of transitions between disease stages.

The problem is formulated as a nine variable bi-objective optimization problem: the aggregated difference for two target quantities. It corresponds to the calibration of the CityCOVID parameters  $\theta$  listed in Table 3, each normalized to  $[0, 1]$ . Model outputs are compared against two empirical data sources obtained through the City of Chicago data portal City of Chicago (2022):  $\mathbf{H}$  the daily census of hospital beds occupied by COVID-19 patients and  $\mathbf{D}$  the COVID-19 attributed death data in and out of hospitals, both for residents of Chicago. We used an exponentially weighted error function  $L(\theta, T_i, \hat{T}_i, d), i \in \{\mathbf{H}, \mathbf{D}\}$ , with daily discount rate  $d$  tuned to 98% and 95% for  $\mathbf{H}$  and  $\mathbf{D}$ , with the corresponding observed (resp. simulated) time-series denoted by  $T$  and  $\hat{T}$  giving the objectives.

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  $\approx 10\text{min}$ , with a low signal-to-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 non-dominated solutions. This data set is available in the supplementary material.

Contrary to the previous test cases that were defined over continuous domains, for testing qHSRI we use this existing data set. The initial design is a random subset of the data: 25,000 simulation results over 2500 unique designs with a degree of replication of 10, out of 50,585 simulations over 5075 unique designs, with a degree of replication between 3 and 10. These correspond to results given by IMABC (akin to a non uniform sampling). 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 to 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 RO.

## 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

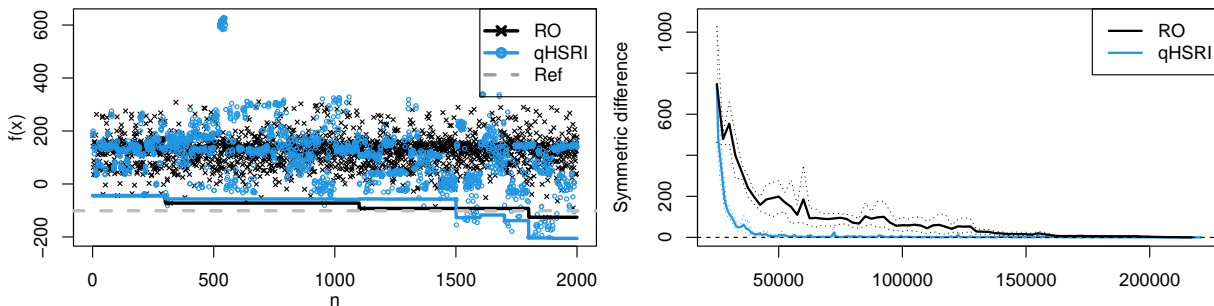


Figure 5: Left: optimality gap for the Lunar lander problem (one single run) with the evaluated values and estimated minimum found. Right: results on CityCOVID data set over 5 macro-runs, thin dashed lines are 5% and 95% quantiles.

of replication. Hence, this proposed approach makes a sensible, simple, yet efficient, baseline for massive batching. Possible extensions could take 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, 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  $\beta$ s, is left for future work.

## 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., Ma, P., Mondal, A., et al. (2022). Analyzing stochastic computer models: A review with opportunities. *Statistical Science*, 37(1):64–89.
- 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.
- Cajka, J. C., Cooley, P. C., and Wheaton, W. D. (2010). Attribute assignment to a synthetic population in support of agent-based disease modeling. *Methods report (RTI Press)*, 19(1009):1–14.
- Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., and Richet, Y. (2014a). 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.
- Chevalier, C., Ginsbourger, D., and Emery, X. (2014b). Corrected kriging update formulae for batch-sequential data assimilation. In *Mathematics of Planet Earth*, pages 119–122. Springer.
- City of Chicago (2022). Data Portal.
- 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.
- Daulton, S., Balandat, M., and Bakshy, E. (2021). Parallel bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *arXiv preprint arXiv:2105.08195*.
- Daxberger, E. A. and Low, B. K. H. (2017). Distributed batch gaussian process optimization. In *International Conference on Machine Learning*, pages 951–960. PMLR.
- De Ath, G., Everson, R. M., Rahat, A. A., and Fieldsend, J. E. (2021). Greed is good: Exploration and exploitation trade-offs in Bayesian optimisation. *ACM Transactions on Evolutionary Learning and Optimization*, 1(1):1–22.
- 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.
- Garnett, R. (2022). Bayesian optimization.
- 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 Nieuwenhuyse, I. (2020). A multiobjective stochastic simulation optimization algorithm. *European Journal of Operational Research*, 284(1):212–226.
- Goubier, T., Rakowsky, N., and Harig, S. (2020). Fast tsunami simulations for a real-time emergency response flow. In *2020 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC)*, pages 21–26.
- 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.
- Gupta, S., Shilton, A., Rana, S., and Venkatesh, S. (2018). Exploiting strategy-space diversity for batch bayesian optimization. In *International conference on artificial intelligence and statistics*, pages 538–547. PMLR.
- 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.
- Macal, C. M., Collier, N. T., Ozik, J., Tatara, E. R., and Murphy, J. T. (2018). ChiSIM: An Agent-Based Simulation Model of Social Interactions in a Large Urban Area. In *2018 Winter Simulation Conference (WSC)*, pages 810–820.
- Mandel, J., Vejmelka, M., Kochanski, A., Farguella, A., Haley, J., Mallia, D., and Hilburn, K. (2019). An interactive data-driven hpc system for forecasting weather, wildland fire, and smoke. In *2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC)*, pages 35–44.
- 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.
- Mutny, M. and Krause, A. (2018). Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 9005–9016. Curran Associates, Inc.
- 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 (2023). *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.



## A Batch Bayesian optimization

### A.1 Synchronous batch BO

The standard batch- (or multi-points) BO loop, where parallel evaluations are conducted in parallel and waiting for all of them to finish, is presented in Algorithm 2.

---

**Algorithm 2** Pseudo-code for batch BO

---

**Input:**  $N_{max}$  (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**
- 

Next we discuss briefly the adaptations required for the asynchronous setup.

### A.2 Asynchronous batch BO

The main feature of the portfolio approach is to give a weight vector corresponding to Pareto optimal points on the mean/standard deviation Pareto front  $\mathcal{P}$ . In the synchronous setting,  $q$  points are selected based on the weights. That is, the  $q$  best ones in the noiseless setting while replicates can occur in the noisy setting. If  $q'$  additional points can be evaluated, without new data available (or waiting for new candidates to evaluate), the change in the asynchronous setup amounts to considering that  $q + q'$  batch points (assets) can be selected according to the weights. The  $q$  first ones will remain unchanged while  $q'$  additional can be run. An example is provided in the Figure 6. We note that for our main motivating example, the CityCOVID simulations, the running times were more or less uniform, hence not benefiting much from a batch asynchronous setup.




Figure 6: Batch allocation on indices according to weights  $w$  (black points) for 10 candidates.  $q = 5$  designs are initially allocated, as represented by the black segments. If  $q' = 2$  additional designs can be evaluated, then the allocation is recomputed with the same weights and leads to the red dashed segments.

## B Additional experimental results

We complement the results of Section 4 on noisy problems with results on noiseless ones. The **R** package **DiceKriging** (Roustant et al., 2012) is used for deterministic GP modeling. In this deterministic setup,  $d$  is increased to accommodate larger batches via repeated versions of these problem as used, e.g., in Oh et al. (2018).

The timing results of all synthetic benchmarks are presented in Tables 1 and 2, with best results in bold. qHSRI or PF are always much faster than the alternatives, even the approximated qEI criterion (qAEI) or qTS. The R code (R Core Team, 2023) of the approach and the CityCOVID data are available as supplementary material. Results have been obtained in parallel on dual-Xeon Skylake SP Silver 4114 @ 2.20GHz (20 cores) and 192 GB RAM (or similar nodes). Lunar lander tests have been run on a laptop.

$f-d-q$	qEI	qAEI	qHSRI	qPF	qTS
B-12-10	24725 (133)	5901 (54)	1103 (104)	<b>1024</b> (160)	1835 (5)
B-12-25	—	5856 (61)	451 (46)	<b>423</b> (66)	1760 (3)
H-12-10	24258 (115)	4957 (131)	<b>969</b> (38)	1266 (144)	1938 (22)
H-12-25	—	4585 (129)	<b>419</b> (24)	510 (76)	1816 (9)
B(h)-2-25	—	2470 (93)	<b>57</b> (3)	342 (36)	305 (46)
H(h)-6-25	—	22012 (717)	<b>288</b> (33)	4714 (702)	5327 (503)

Table 1: Averaged timing results (in s), with standard deviation in parenthesis. B is for Branin, H for Hartmann6, (h) for heteroscedastically noisy problems and '—' indicates when non applicable.

$f-d-q$	qEHI	qHSRI	qPF
P1-6-50	16308 (2118)	678 (25)	<b>580</b> (26)
P1(h)-2-50	35350 (1363)	<b>567</b> (59)	5110 (407)
P2-6-50	14022 (1818)	635 (24)	<b>553</b> (20)
P2(h)-2-50	37386 (1123)	<b>608</b> (44)	5148 (414)

Table 2: Averaged timing results (in s), with standard deviation in parenthesis. (h) is for heteroscedastically noisy problems.

The detailed results of the performance over time are given in Figure 7, where qHSRI show better results than the alternatives, closely matched by PF except on Branin where it does worse. qTS performs badly as the discretization is insufficient given the problem dimension.

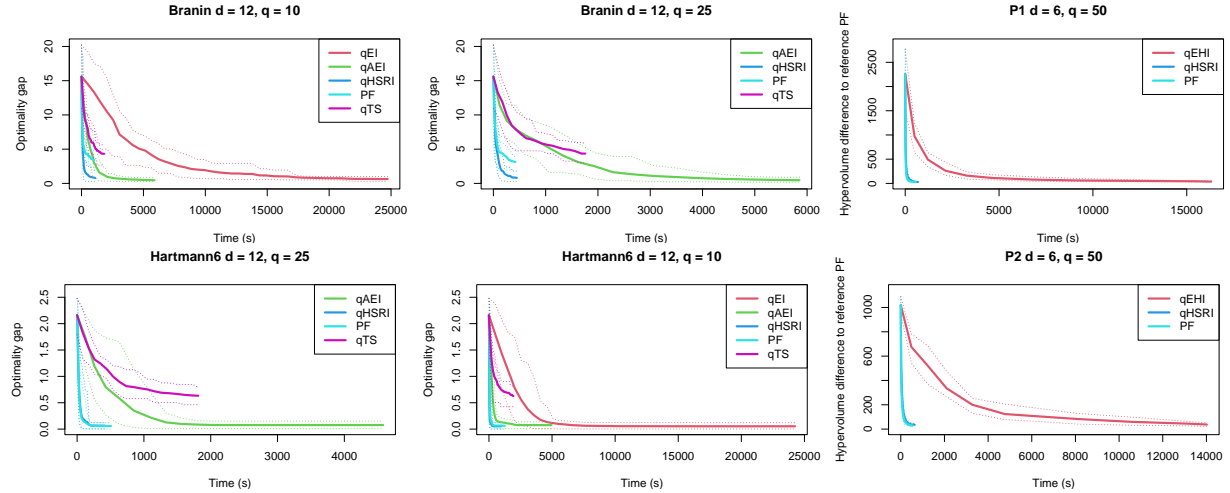


Figure 7: Noiseless results over time. Optimality gap or hypervolume difference over 20 macro-runs is given. Thin dashed lines are 5% and 95% quantiles.

The results over iterations of Figure 8 shows that the performance of qHSRI matches the one of qEI and improves over qEHI. qTS underperforms but it is still better than RO. PF shows similar results to qHSRI on the multi-objective problems and Hartmann6 with  $q = 25$ , but is worse in the other test cases. For the smallest batch size,  $q = 10$ , qAEI matches the performance of qEI, while being faster.

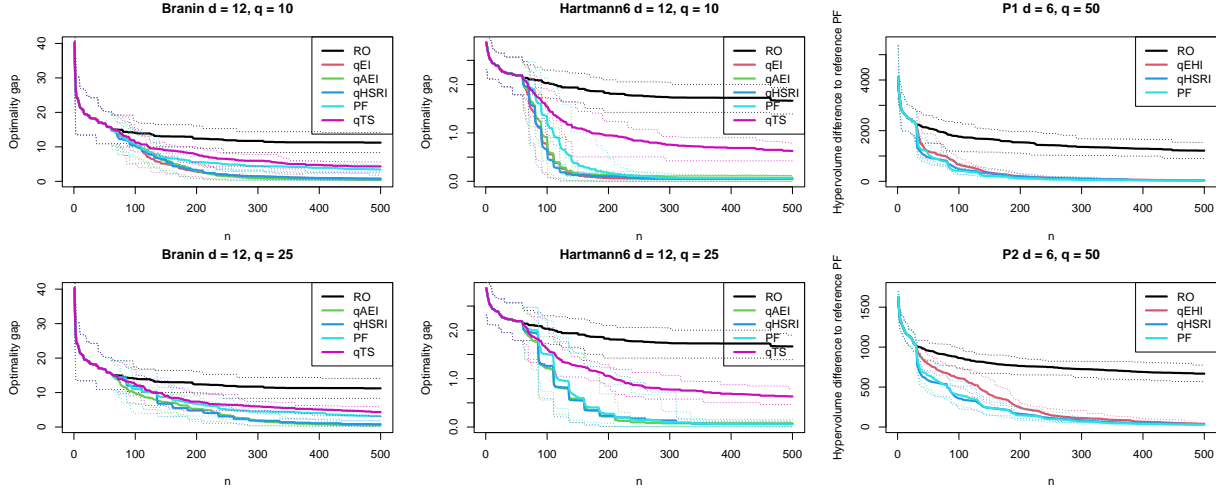


Figure 8: Noiseless results over iterations. Optimality gap or hypervolume difference over 20 macro-runs is given. Thin dashed lines are 5% and 95% quantiles.

To highlight that replication occurs, Figure 9 represents the number of unique design over time for the mono-objective noisy problems. Notice that replication is present in the initial design and that it can happen for all methods when sampling vertices of the hypercube.




Figure 9: Number of unique designs over iterations for the mono-objective test problems. Thin dashed lines are 5% and 95% quantiles.

## C Details on CityCOVID calibration parameters

The nine variables of the simulator are given in Table 3.

$\theta$	$\pi(\theta)$	Description
$\theta_1$	$U(60, 190)$	Initial number of exposed (infected but not infectious) agents at the beginning of the simulation
$\theta_2$	$U(0.03, 0.1)$	Base hourly probability of transmission between one infectious and one susceptible person occupying the same location
$\theta_3$	$U(0, 0.3)$	Magnitude of seasonality effect
$\theta_4$	$U(0.5, 1)$	Per person probability of infection scaling factor due to ratio of infectious versus susceptible people in a location
$\theta_5$	$U(0.2, 0.7)$	Effective infectivity during isolation in household
$\theta_6$	$U(0.1, 0.7)$	Effective infectivity during isolation in nursing home
$\theta_7$	$U(300, 700)$	Simulation time (hrs) corresponding to March 27, 2020
$\theta_8$	$U(0.1, 0.8)$	Reduction in out of household activities
$\theta_9$	$U(0.01, 0.3)$	Reduction in transmission due to individual protective behaviors

Table 3: CityCOVID calibration parameters  $\theta$  and priors  $\pi(\theta)$ .