

Reduced-Space Multi-Fidelity Bayesian Optimization of Process Simulation Models

Niki Triantafyllou, Andrea Bernardi, and Maria M. Papathanasiou*

Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, Imperial College London, UK
{niki.triantafyllou20, a.bernardi13, maria.papathanasiou11}@imperial.ac.uk

Abstract. Optimizing industrial process flowsheets is often computationally prohibitive due to the high cost of rigorous simulations and the curse of dimensionality inherent in complex design spaces. To address these challenges, we present a reduced-space multi-fidelity Bayesian optimization (RS-MFBO) framework designed for high-dimensional, expensive black-box functions. The approach integrates Global Sensitivity Analysis (GSA) for dimensionality reduction with a fidelity-augmented Gaussian process that captures correlations between low-cost approximations and expensive high-fidelity evaluations via a linear-truncated kernel. A cost-aware acquisition strategy, augmented with cooldown and promotion mechanisms, adaptively guides the allocation of samples across fidelities. The framework is validated on two distinct industrial process simulators: a plasmid DNA bioprocess in *SuperPro Designer* and a green fuel synthesis plant in *Aspen HYSYS*. Results across diverse economic and physical objectives demonstrate that the proposed method substantially reduces the number of high-fidelity simulator evaluations while maintaining competitive optimization performance compared to single-fidelity baselines. These results highlight RS-MFBO as a scalable, simulator-agnostic approach for cost-constrained black-box optimization.

Keywords: Multi-fidelity Bayesian Optimization · Bayesian Optimization · Gaussian Processes · Active Learning.

1 Introduction

Systematic process design and optimization aim to determine operating and design conditions that maximize process performance (e.g., throughput, yield, energy efficiency) or minimize cost and environmental impact. (Bio)chemical *process simulators* such as Aspen and SuperPro Designer are widely used for these problems owing to their extensive model libraries, integrated physical-property databases, and user-friendly interfaces. While such simulators provide high-fidelity (HI) representations of process behavior, their governing equations are not exposed, so derivatives are unavailable and gradient-based methods are

* Corresponding author.




Fig. 1. High-fidelity black-box simulator models: (a) **pDNA bioprocess** (SuperPro Designer), (b) **green fuel synthesis** (Aspen HYSYS) [24].

inapplicable. Consequently, sequential-modular flowsheet optimization relies on derivative-free, sample-efficient strategies.

Bayesian optimization (BO) offers a principled framework for optimizing expensive black-box functions under limited evaluations [8, 1, 16, 23, 9]. However, in flowsheet optimization settings, BO faces three key challenges: (i) high-dimensional design spaces; (ii) costly and occasionally non-convergent simulator evaluations (necessitating random restarts or discarding failed runs); and (iii) operational constraints that restrict feasible operating regions. Constrained BO is therefore required to keep exploration within known safe limits [3]. At the same time, single-fidelity simulation-based BO is inefficient when simulator calls are expensive and unreliable, while surrogate-only BO tends to lose accuracy when extrapolating beyond its training domain [23].

Multi-fidelity Bayesian optimization (MFBO) mitigates the challenges of costly simulator evaluations and imperfect surrogate models by combining multi-

ple sources of information, typically a cheap low-fidelity (LO) approximation and an expensive high-fidelity (HI) simulator, within a unified probabilistic multi-fidelity model [11, 12, 5, 4]. By jointly learning cross-fidelity correlations and using cost-aware acquisition functions that balance exploration, exploitation, and fidelity accuracy, MFBO efficiently allocates evaluations under a fixed computational budget.

Even with MFBO, high dimensionality remains a bottleneck in flowsheet optimization. To address this, *global sensitivity analysis (GSA)* can be employed to eliminate low-influence variables prior to optimization. GSA provides an interpretable, model-agnostic means of quantifying the contribution and interaction effects of each input on the objective function, enabling an explainable form of dimensionality reduction.

Building on our earlier work on reduced-space single-fidelity BO [23], we now introduce a *reduced-space multi-fidelity Bayesian optimization* (RS-MFBO) framework that couples surrogate and simulator evaluations through a fidelity-augmented Gaussian process with a linear-truncated cross-fidelity kernel. A cost-aware acquisition policy incorporating cooldown and promotion mechanisms adaptively balances LO and HI queries throughout the search. We assess the algorithmic robustness of this framework against two benchmark case studies utilizing different (bio)chemical process simulators: (a) plasmid DNA production in *SuperPro Designer* with 18 mixed continuous and discrete decision variables, and (b) green fuel (dimethyl ether) production in *Aspen HYSYS* with 14 mixed continuous and discrete decision variables. By benchmarking against single-fidelity and surrogate-based baselines across twelve diverse objectives (six per case study), we show that RS-MFBO achieves state-of-the-art performance with significantly reduced computational overhead (Figure 1).

2 Related Work

High-Dimensional Bayesian Optimization. Scaling BO to high-dimensional design spaces ($D > 20$) is a longstanding challenge. A common strategy involves embedding the high-dimensional problem into a lower-dimensional latent subspace. This includes random embedding methods like REMBO [25] and HeSBO [15], as well as recent extensions for discrete search spaces that project integer variables into continuous latent manifolds [9]. Other approaches, such as TuRBO [7], manage dimensionality by restricting the search to local trust regions. While numerically effective, these embedding and projection methods operate in latent spaces that lack physical meaning. In process systems engineering, maintaining interpretability is crucial. Decision-makers need to understand which specific physical variables (e.g., temperatures, flow rates) drive performance. This motivates our use of Global Sensitivity Analysis (GSA) [18] to identify the active subspace directly in the native coordinate system, preserving engineering insight.

BO in Process Systems Engineering. Bayesian optimization is gaining traction in the chemical engineering community as a sample-efficient alternative

to derivative-free solvers. Recent applications cover a broad spectrum of process systems problems, including reactor design [19], bioprocess development [13], and molecular discovery [20, 14]. Specialized frameworks like ENTMOOT [22] have also been developed to handle tree-based surrogate models within BO. Although some of these studies have successfully introduced multi-fidelity strategies to specific domains, such hierarchical data structures are often neglected in broader process optimization tasks, where single-fidelity surrogates remain the standard.

Multi-Fidelity Acquisition Strategies. Multi-fidelity BO (MFBO) integrates information sources of varying costs and accuracies. Existing methods often assume a discrete fidelity hierarchy (e.g., MISO [17]) or rely on information-theoretic acquisition functions, such as Multi-Fidelity Entropy Search (MUMBO) [21] or Knowledge Gradient (MF-KG) [26]. While these advanced acquisition functions offer rigorous theoretical bounds, they incur significant computational overhead due to the need for numerical integration or inner optimization loops. In our context, where high-fidelity simulations require only 5–45 seconds, the wall-time cost of optimizing an expensive acquisition function (like KG) can exceed the cost of the simulation itself, negating the benefits of MFBO. To ensure the framework remains practical for intermediate-cost simulators, we adopt the continuous fidelity relaxation proposed by Wu et al. [26], but pair it with a lightweight, analytical cost-aware UCB strategy instead of the computationally expensive Knowledge Gradient.

3 Methodology

3.1 Problem Setting

The optimization task seeks to identify process design and operating variables $\mathbf{x} \in \mathcal{X}$ that optimize a high-fidelity (HI) simulator objective. Formally, the problem is expressed as a constrained black-box optimization:

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}_{\text{feas}}} \quad & f_{\text{HI}}(\mathbf{x}) \\ \text{s.t.} \quad & g_j(\mathbf{x}) \leq 0, \quad j = 1, \dots, m, \end{aligned} \tag{1}$$

where $\mathcal{X}_{\text{feas}} \subseteq \mathcal{X}$ denotes the feasible input domain satisfying all operational and physical constraints $g_j(\mathbf{x})$. Evaluations of f_{HI} are computationally expensive and derivative-free, motivating the use of a low-fidelity surrogate f_{LO} that approximates f_{HI} at reduced cost. The proposed multi-fidelity Bayesian optimization (MFBO) framework (Algorithm 1) adaptively allocates evaluations between f_{LO} and f_{HI} to efficiently solve (1) under a finite computational budget.

3.2 Input-Augmented Multi-Fidelity Gaussian Process

We model both fidelities jointly as a single Gaussian process (GP) defined over an augmented input space $(\mathbf{x}, s)$, where $s \in [0, 1]$ represents the fidelity level ($s=0$ for low fidelity and $s=1$ for high fidelity). This continuous formulation

captures smooth correlations between fidelities while avoiding the need for separate surrogate models.

The prior over the latent function $f(\mathbf{x}, s)$ is

$$f(\mathbf{x}, s) \sim \mathcal{GP}(m(\mathbf{x}, s), k((\mathbf{x}, s), (\mathbf{x}', s'))), \quad (2)$$

with a constant mean $m(\mathbf{x}, s) = \mu_0$ and a separable covariance:

$$k((\mathbf{x}, s), (\mathbf{x}', s')) = k_x(\mathbf{x}, \mathbf{x}') k_s(s, s'), \quad k_s(s, s') = \min(s, s'). \quad (3)$$

Here, k_x is a Matérn-5/2 kernel applied to normalized inputs with standardized outputs. The fidelity kernel $k_s(s, s') = \min(s, s')$ is a *linear-truncated kernel* that encodes a hierarchical dependence: low-fidelity information informs the high-fidelity surface, but not conversely.

Conditioning on observed data $\mathcal{D} = \{(\mathbf{x}_i, s_i, y_i)\}_{i=1}^N$ yields the posterior predictive distribution

$$p(f_* | \mathbf{x}_*, s_*, \mathcal{D}) = \mathcal{N}(\mu_*(\mathbf{x}_*, s_*), \sigma_*^2(\mathbf{x}_*, s_*)), \quad (4)$$

with standard GP mean and variance updates. This input-augmented formulation is mathematically equivalent to the classical autoregressive multi-fidelity model [12], where $f_{\text{HI}}(\mathbf{x}) = \rho f_{\text{LO}}(\mathbf{x}) + \delta(\mathbf{x})$ and $\delta(\mathbf{x})$ represents the discrepancy between fidelities. By coupling both fidelities in a single joint prior, the model enables uncertainty reduction in the high-fidelity predictions even when only low-fidelity evaluations are added. This provides a compact, numerically stable alternative to multi-output co-kriging for computationally expensive process simulations.

3.3 Cost-Aware Acquisition with Cooldown and Promotion

At iteration t , we optimize the upper-confidence-bound (UCB) acquisition separately at each fidelity and score them cost-adjustedly:

$$\text{UCB}_s(\mathbf{x}) = \mu_s(\mathbf{x}) + \beta \sigma_s(\mathbf{x}), \quad s \in \{\text{LO}, \text{HI}\}, \quad (5)$$

$$J_s = \frac{\text{UCB}_s(\mathbf{x})}{c_s^\alpha}. \quad (6)$$

The fidelity with the larger J_s is selected, but a high-fidelity (HI) evaluation is *forced* every H steps to prevent over-reliance on the low-fidelity (LO) surrogate (cooldown). Every P steps, the top- K LO points with the highest predicted improvement are *promoted* for HI reevaluation.

Before promotion, candidate LO points are checked for redundancy against existing HI evaluations. We remove any point $\mathbf{x}_{\text{new}}$ that lies within a small ℓ_∞ -norm distance of an existing HI point $\mathbf{x}_{\text{old}}$:

$$\|\mathbf{x}_{\text{new}} - \mathbf{x}_{\text{old}}\|_\infty = \max_i |x_{\text{new},i} - x_{\text{old},i}| < \varepsilon,$$

Algorithm 1 Reduced-space cost-aware multi-fidelity Bayesian optimization

Require: Feasible domain $\mathcal{X}_{\text{feas}}$, budget B , fidelity costs $c_{\text{LO}} < c_{\text{HI}}$, UCB parameters (β, α) , cooldown period H , promotion interval P , and top- K candidates

- 1: **GSA screening:** compute total-order Sobol' indices S_i^T and retain dominant variables to form reduced space $\mathcal{X}'_{\text{feas}}$
- 2: **Initialization:** sample Sobol' points at both fidelities in $\mathcal{X}'_{\text{feas}}$; evaluate f_{LO} and f_{HI} ; set total cost C
- 3: **while** $C < B$ **do**
- 4: Fit MF-GP
- 5: Maximize $\text{UCB}_s(\mathbf{x}) = \mu_s(\mathbf{x}) + \beta\sigma_s(\mathbf{x})$ for $s \in \{\text{LO}, \text{HI}\}$
- 6: Compute cost-adjusted scores $J_s = \text{UCB}_s/c_s^\alpha$
- 7: **if** cooldown reached or $J_{\text{HI}} \geq J_{\text{LO}}$ **then**
- 8: Evaluate $f_{\text{HI}}(\mathbf{x}_{\text{HI}})$
- 9: **else**
- 10: Evaluate $f_{\text{LO}}(\mathbf{x}_{\text{LO}})$
- 11: **end if**
- 12: **if** iteration mod $P = 0$ **then**
- 13: Promote top- K LO points to HI and re-evaluate (duplicates removed)
- 14: **end if**
- 15: **end while**
- 16: **return** best observed high-fidelity objective

with $\varepsilon = 10^{-6}$ in normalized coordinates. This deduplication avoids redundant simulator calls and numerical instability due to repeated evaluations at nearly identical conditions. The surviving points are then evaluated at high fidelity, and their results are incorporated into the dataset used for subsequent model updates. The total computational budget is updated accordingly.

The MFBO loop proceeds until the total computational budget is exhausted or no further improvement in the best observed high-fidelity objective is detected over a fixed number of iterations, indicating convergence.

3.4 Reduced Space via Variance-Based Global Sensitivity Analysis

To alleviate the curse of dimensionality and improve model conditioning, we perform a variance-based *global sensitivity analysis* (GSA) using total-order Sobol' indices [18]. Let $f(\mathbf{x})$ denote the model output and $\text{Var}(f)$ its total variance over the input domain $\mathcal{X}$. For each input variable x_i , the total-order index S_i^T quantifies its overall contribution (main effect and all interactions) to the output variance:

$$S_i^T = 1 - \frac{\text{Var}_{\mathbf{x}_{\sim i}}(\mathbb{E}[f(\mathbf{x}) \mid \mathbf{x}_{\sim i}])}{\text{Var}(f)}, \quad (7)$$

where $\mathbf{x}_{\sim i}$ denotes all inputs except x_i . Variables with negligible S_i^T have limited influence on the objective and can be safely fixed at nominal values.

We retain the subset of variables $\mathcal{I} = \{i : S_i^T \geq \tau\}$ that collectively explain a fraction τ of the total output variance of f .

4 Experimental Setup

4.1 Case Studies

To validate the generalizability of RS-MFBO, we apply the framework to two distinct industrial process simulators (Figure 1) representing different problem classes.

Case Study I: Biopharmaceutical Process (SuperPro Designer). We optimize a plasmid DNA (pDNA) production process [24]. The system is modeled as a steady-state representation of a batch process in *SuperPro Designer*. This setup captures the aggregated material and energy balances of a complete production cycle, involving strictly sequential unit operations (fermentation, lysis, chromatography).

- **Search Space:** The problem involves $D = 18$ mixed continuous and discrete variables (e.g., equipment sizing, batch scheduling parameters, flow rates).
- **Objectives:** We optimize six KPIs: (1) *Production Cost*, (2) *Operating Expenditure (OpEx)*, (3) *Capital Expenditure (CapEx)*, (4) *Batch Size*, (5) *Batch Time*, and (6) *Cycle Time*. This set captures the trade-offs between cost minimization and throughput maximization under strict scheduling constraints. Two linear input constraints enforce the increasing duration of fermentation stages:

$$t_{\text{flask}} \leq t_{\text{seed}} \leq t_{\text{main}}$$

Case Study II: Chemical Synthesis Process (Aspen HYSYS). We optimize a green fuel synthesis plant producing Dimethyl Ether (DME) from captured CO_2 and green H_2 . Modeled in *Aspen HYSYS*, this is a steady-state continuous process characterized by complex thermodynamic interactions and rigorous recycle loops, which frequently induce simulator non-convergence.

- **Search Space:** The problem involves $D = 14$ mixed continuous and discrete variables (e.g., reactor pressures, column tray counts, recycle ratios).
- **Objectives:** We optimize six KPIs including: (1) *DME Production Rate*, (2) *Energy Efficiency*, (3) *Carbon Efficiency*, (4) *OpEx*, (5) *CapEx*, and (6) *Production Cost*. This selection tests the algorithm’s ability to handle highly non-convex thermodynamic landscapes.

4.2 Baselines

We compare the proposed RS-MFBO against five baselines:

1. **Sobol’ Sampling:** Quasi-random search (performance lower bound).
2. **BO (Standard):** Vanilla high-fidelity Bayesian Optimization on the full input space.
3. **RS-BO:** Single-fidelity BO operating on the reduced subspace defined by GSA.

4. **RS-ANN-BO:** A surrogate-based approach where BO optimizes the low-fidelity ANN on the reduced subspace defined by GSA, validated once at the end.
5. **RS-ANN-MILP:** Deterministic global optimization of the ReLU ANN surrogate on the reduced subspace defined by GSA via Mixed-Integer Linear Programming [6].

4.3 Implementation Details

A VBA-Python COM interface was developed to enable automated communication with the SuperPro Designer and Aspen HYSYS simulators. Experiments were conducted using a fixed total cost budget. High-fidelity simulator evaluations were assigned a normalized cost $c_{\text{HI}} = 10.0$, while low-fidelity surrogate evaluations were assigned $c_{\text{LO}} \approx 10^{-1}$. GSA screening was performed using the SALib library [10] using a total-order Sobol’ index threshold of $\tau = 0.01$. The Multi-Fidelity GPs were implemented in BoTorch [2] using a linear-truncated kernel structure. UCB parameters were set to $\beta = 15.0$ and $\alpha = 0.1$, with cooldown period $H = 10$, promotion interval $P = 10$, and $K = 1$ candidate promoted per interval. Each run was initialised with $N_{\text{HI}} = 2$ and $N_{\text{LO}} = 2$ Sobol’ points per fidelity. All results are reported as the median and interquartile range (IQR) over 10 random seeds. Neural network training was GPU-accelerated, while Bayesian optimization and MILP were executed on CPU.

Apart from the pure Sobol’ sampling baseline, all methods are initialized with the best-performing design points identified by the GSA-based screening procedure, whereas the vanilla BO baseline starts from independent Sobol’ samples.

5 Results and Discussion

The proposed reduced-space MFBO framework was evaluated against five single- and surrogate-based optimization baselines across two distinct industrial process case studies: a batch plasmid DNA (pDNA) production process modeled in SuperPro Designer and a continuous green fuel (dimethyl ether) synthesis plant modeled in Aspen HYSYS. In total, twelve diverse key performance indicators (six per case study) were optimized, covering economic metrics like capital and operating expenditures, as well as physical performance indicators such as batch time, cycle time, and carbon efficiency. Each objective was optimized independently under identical initial designs and total cost budgets to ensure a fair comparison of sample efficiency across methods. Critical variables and low-fidelity ANN surrogates were obtained via GSA screening as described in the methodology.

5.1 Overall optimization performance

The reduced-space multi-fidelity BO (RS-MFBO) demonstrates robust convergence across all objectives in both case studies, despite the distinct nature of their optimization landscapes.




Fig. 2. Optimization results for all objectives: best observed high-fidelity value vs. function evaluations for Case Study I. Surrogate-based methods are validated with the simulator.

In the pDNA case study, the problem involves strictly sequential unit operations and scheduling constraints such as the non-decreasing fermentation times which are enforced across all benchmark methods ($t_{\text{flask}} \leq t_{\text{seed}} \leq t_{\text{main}}$).

Optimization trajectories for the pDNA bioprocess are presented in Figure 2. RS-MFBO demonstrates robust convergence across most objectives, though the results highlight the inherent trade-offs of dimensionality reduction. For economic objectives such as Production Cost (Fig. 2f), RS-MFBO matches the final median performance of the high-fidelity baselines, converging to the optimum within the first 10 to 15 evaluations. Notably, for CapEx (Fig. 2d), RS-MFBO achieves a visibly lower final median value than both vanilla BO and reduced-space BO (RS-BO), suggesting that the multi-fidelity exploration helped escape local optima that trapped the greedy single-fidelity methods. However, in the case of OpEx (Fig. 2e), the reduced-space methods (both RS-BO and RS-MFBO) slightly underperform compared to the full-space vanilla BO. This behavior is expected in reduced-space optimization; by fixing non-dominant variables to nominal values, the optimizer operates on a restricted manifold that may exclude the true global optimum if those non-critical variables have small but non-zero effects.

Performance on operational KPIs varies by difficulty. For Batch Size (Fig. 2a), all optimization-based methods converge to a similar high-throughput solution, significantly outperforming random sampling. In contrast, Cycle Time (Fig. 2c) proves to be an easy optimization landscape where even the baseline Sobol sampling identifies the optimal solution immediately. In these simpler tasks, RS-MFBO still offers value by confirming optimality with fewer expensive checks




Fig. 3. Optimization results for all objectives for case study II: best observed high-fidelity value vs. function evaluations. Surrogate-based methods are validated with the simulator.

than standard BO. Throughout these experiments, the pure surrogate-based baselines (RS-ANN-BO and RS-ANN-MILP) appear as single points in the figures. While these methods incur the lowest computational cost, their accuracy is unreliable. RS-ANN-MILP often identifies designs that appear optimal *in silico* but degrade significantly when validated against the true high-fidelity simulator, highlighting the necessity of the feedback loop provided by RS-MFBO.

The optimization results for the continuous DME synthesis process are summarized in Figure 3. This case study poses a distinct challenge due to complex thermodynamic recycle loops that frequently destabilize the process simulation. RS-MFBO leverages the smooth low-fidelity surrogate to navigate this landscape efficiently. This is most evident in the *Energy Efficiency* maximization (Fig. 3b). Here, RS-MFBO exploits the high-quality region identified by the low-fidelity surrogate to initialize the search near 80% efficiency, avoiding the initial exploration phase that delays standard BO and RS-BO, which start significantly lower (between 70% and 76%). While the single-fidelity methods eventually climb to competitive values, RS-MFBO maintains a performance lead throughout the budget, reaching peak efficiencies of approximately 84%.

Similar advantages are observed for the *DME Production Rate* (Fig. 3a), where RS-MFBO identifies the optimal operating region (approx. 4.0 kg/s) almost immediately. In contrast, standard vanilla BO shows a delayed convergence, requiring over 40 function evaluations to match the performance that RS-MFBO achieves in fewer than 10. For economic objectives like *OpEx* (Fig. 3d) and *CapEx* (Fig. 3e), the method exploits the high quality of the initial surrogate to flatten the cost curve instantly. While vanilla BO eventually converges to a similar




Fig. 4. High fidelity (HI) evaluations vs. iteration for the pDNA case study.

cost, it incurs a significant regret penalty during the initial exploration phase. The difference in convergence speed between RS-MFBO and the single-fidelity baselines across these objectives highlights the method’s ability to handle high-dimensional search spaces effectively.

5.2 Effect of reduced-space modeling

The comparison between vanilla Bayesian optimization (BO), reduced-space BO (RS-BO), and RS-MFBO highlights the benefits of GSA-based dimensionality reduction across both simulators. In several SuperPro objectives, RS-BO converges faster than vanilla BO, indicating that restricting the search to the most influential variables can improve sample efficiency. RS-MFBO inherits this advantage and further augments it with a fidelity-aware GP and cost-aware acquisition policy, yielding the best overall cost-performance trade-off. From a modeling standpoint, operating in the reduced space improves GP conditioning and leads to more stable hyperparameter estimates. In the full 18-dimensional pDNA problem and the 14-dimensional DME problem, kernel hyperparameter optimization is often ill-conditioned, leading to over-smoothing and large posterior variance. By retaining only the most influential variables (typically 6–8 depending on the objective), the reduced-space formulation produced more stable marginal likelihood fits and better-calibrated uncertainty estimates. Furthermore, the reduced space significantly accelerates the internal acquisition optimization step, cutting its wall-time by approximately 40–60% on average compared to the full-space operation.




Fig. 5. High fidelity (HI) evaluations vs. iteration for the DME case study.

5.3 Acquisition dynamics and use of high-fidelity evaluations

Analysis of the cumulative high-fidelity (HI) simulator evaluations reveals a characteristic “staircase” usage profile for RS-MFBO across all 12 objectives (see Figures 4, 5). The flat regions of this profile correspond to extended phases of low-fidelity exploration where the acquisition function exploits the cheap surrogate or explores regions with high surrogate uncertainty. Vertical jumps occur only when the cooldown mechanism forces a correction or when the promotion heuristic identifies a promising candidate from the low-fidelity model. This behavior arises naturally from the cost-aware acquisition policy, which prioritizes the cheaper ANN surrogate unless the cost-adjusted UCB score or the cooldown interval indicates the need for high-fidelity refinement.

Across both the SuperPro and Aspen HYSYS case studies, RS-MFBO uses substantially fewer HI simulator calls than the single-fidelity BO baselines. While vanilla BO and reduced-space single-fidelity BO query the simulator at every iteration, RS-MFBO performs the majority of its evaluations on the low-fidelity surrogate and relies on selective high-fidelity queries to maintain surrogate accuracy. The ℓ_∞ -based deduplication step further avoids redundant re-evaluations at nearly identical conditions. Overall, RS-MFBO attains competitive or improved final objective values while using approximately 65–80% fewer high-fidelity evaluations than pure BO under the same total cost budget. This efficiency is consistent across both case studies.

5.4 Computational time

Figure 6 compares the wall-time for all methods. As expected, pure BO incurs the highest computational cost because every step involves a simulator call in




Fig. 6. Execution time by method and objective.

the full input space. In the SuperPro Designer case study (Fig. 6a), individual simulations take approximately 10–15 seconds. However, the total runtime for vanilla BO is disproportionately high. This is primarily because optimizing the acquisition function within the full 18-dimensional space, subject to the two constraints, is computationally intensive. Reduced-space single-fidelity BO (RS-BO) reduces wall-time relative to BO by operating in a smaller subspace, but its reliance on high-fidelity evaluations still dominates the overall runtime.

In the Aspen HYSYS case study (Fig. 6b), simulation times vary significantly (ranging from 5 to 45 seconds) depending on the difficulty of converging the recycle loops. Here, RS-MFBO consistently lies among the fastest methods. The wall-time reduction stems from two factors: (i) most iterations query the cheap low-fidelity surrogate, avoiding the heavy overhead of unstable simulator calls, and (ii) GSA restricts both the GP fit and the acquisition optimization to a reduced number of variables. In many objectives, RS-MFBO achieves near-best objective values within a fraction of the runtime of BO and RS-BO, making it more suitable for time- and budget-constrained simulation models. The surrogate-only methods (RS-ANN-BO and RS-ANN-MILP) are the fastest in absolute terms because they avoid simulator calls altogether (except for final validation). However, their performance is ultimately bounded by the accuracy of the ANN. When the surrogate misrepresents the true global optimum, the surrogate-only methods converge to suboptimal designs. In contrast, RS-MFBO offers a better trade-off between computational cost and solution quality by systematically combining low- and high-fidelity information.

6 Conclusion

We presented a reduced-space multi-fidelity Bayesian optimization framework tailored to simulation-based models and evaluated it on two different simulators. The combination of GSA-based dimensionality reduction, a fidelity-augmented Gaussian process with a linear-truncated kernel, and a cost-aware acquisition strategy with cooldown and promotion enabled efficient exploration while limiting the number of expensive simulator calls. Across the KPIs studied, the

method achieved competitive final performance while using substantially fewer HI evaluations than single-fidelity Bayesian optimization.

Future work includes extending the fidelity hierarchy beyond two levels, and developing learned fidelity selection policies for broader applicability.

Acknowledgments. NT and MMP would like to thank Ben Lyons, Benoit Chachuat and Cleo Kontoravdi for their contributions and input on the construction of the flow-sheets and insights on the case studies. NT is thankful for the Marit Mohn Scholarship awarded by the Department of Chemical Engineering, Imperial College London.

References

1. Archetti, F., Candelieri, A.: Bayesian optimization and data science, vol. 849. Springer (2019)
2. Balandat, M., Karrer, B., Jiang, D.R., Daulton, S., Letham, B., Wilson, A.G., Bakshy, E.: BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In: Advances in Neural Information Processing Systems 33 (2020)
3. Candelieri, A.: Sequential model based optimization of partially defined functions under unknown constraints. *Journal of Global Optimization* **79**(2), 281–303 (2021)
4. Candelieri, A., Ponti, A., Archetti, F.: Fair and green hyperparameter optimization via multi-objective and multiple information source bayesian optimization. *Machine Learning* **113**(5), 2701–2731 (2024)
5. Candelieri, A., Ponti, A., Archetti, F., Sabatella, A.: Multiple Information Source Bayesian Optimization. Springer (2025)
6. Ceccon, F., Jalving, J., Haddad, J., Thebelt, A., Tsay, C., Laird, C.D., Misener, R.: Omlt: optimization & machine learning toolkit. *J. Mach. Learn. Res.* **23**(1) (Jan 2022)
7. Eriksson, D., Pearce, M., Gardner, J., Turner, R.D., Poloczek, M.: Scalable global optimization via local bayesian optimization. *Advances in neural information processing systems* **32** (2019)
8. Garnett, R.: Bayesian optimization. Cambridge University Press (2023)
9. Hernández-Morales, G., Cansino-Loeza, B., Jiménez-Gutiérrez, A., Zavala, V.M.: Simulation-based optimization over discrete spaces using projection to continuous latent spaces (2025), <https://arxiv.org/abs/2510.14206>
10. Iwanaga, T., Usher, W., Herman, J.: Toward salib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses. *Socio-Environmental Systems Modelling* **4**, 18155 (May 2022). <https://doi.org/10.18174/sesmo.18155>
11. Kandasamy, K., Dasarthy, G., Schneider, J., Póczos, B.: Multi-fidelity bayesian optimisation with continuous approximations. In: 34th International Conference on Machine Learning, ICML 2017. pp. 2861–2878. International Machine Learning Society (IMLS) (2017)
12. Kennedy, M.C., O’Hagan, A.: Predicting the output from a complex computer code when fast approximations are available. *Biometrika* **87**(1), 1–13 (2000)
13. Martens, A., Neufang, M., Butté, A., von Stosch, M., del Rio Chanona, A., Helleskes, L.M.: Holistic bioprocess development across scales using multi-fidelity batch bayesian optimization (2025), <https://arxiv.org/abs/2508.10970>

14. McDonald, M.A., Koscher, B.A., Canty, R.B., Zhang, J., Ning, A., Jensen, K.F.: Bayesian optimization over multiple experimental fidelities accelerates automated discovery of drug molecules. *ACS central science* **11**(2), 346–356 (2025)
15. Nayebi, A., Munteanu, A., Poloczek, M.: A framework for bayesian optimization in embedded subspaces. In: *International Conference on Machine Learning*. pp. 4752–4761. PMLR (2019)
16. Paulson, J.A., Tsay, C.: Bayesian optimization as a flexible and efficient design framework for sustainable process systems. *Current Opinion in Green and Sustainable Chemistry* **51**, 100983 (2025)
17. Poloczek, M., Wang, J., Frazier, P.: Multi-information source optimization. *Advances in neural information processing systems* **30** (2017)
18. Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., Tarantola, S.: Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer Physics Communications* **181**(2), 259–270 (2010)
19. Savage, T., Basha, N., McDonough, J., Krassowski, J., Matar, O., del Rio Chanona, E.A.: Machine learning-assisted discovery of flow reactor designs. *Nature Chemical Engineering* **1**(8), 522–531 (2024)
20. Shields, B.J., Stevens, J., Li, J., Parasram, M., Damani, F., Alvarado, J.I.M., Janey, J.M., Adams, R.P., Doyle, A.G.: Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **590**(7844), 89–96 (2021)
21. Song, J., Chen, Y., Yue, Y.: A general framework for multi-fidelity bayesian optimization with gaussian processes. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. pp. 3158–3167. PMLR (2019)
22. Thebelt, A., Tsay, C., Lee, R., Sudermann-Merx, N., Walz, D., Shafei, B., Misener, R.: Tree ensemble kernels for bayesian optimization with known constraints over mixed-feature spaces. *Advances in Neural Information Processing Systems* **35**, 37401–37415 (2022)
23. Triantafyllou, N., Lyons, B., Bernardi, A., Chachuat, B., Kontoravdi, C., Papathanasiou, M.M.: Comparative assessment of simulation-based and surrogate-based approaches to flowsheet optimization using dimensionality reduction. *Computers & Chemical Engineering* **189**, 108807 (2024)
24. Triantafyllou, N., Sarkis, M., Krassakopoulou, A., Shah, N., Papathanasiou, M.M., Kontoravdi, C.: Uncertainty quantification for gene delivery methods: A roadmap for pdna manufacturing from phase i clinical trials to commercialization. *Biotechnology Journal* **19**(1), 2300103 (2024)
25. Wang, Z., Hutter, F., Zoghi, M., Matheson, D., De Freitas, N.: Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research* **55**, 361–387 (2016)
26. Wu, J., Toscano-Palmerin, S., Frazier, P.I., Wilson, A.G.: Practical multi-fidelity bayesian optimization for hyperparameter tuning. In: *Uncertainty in Artificial Intelligence*. pp. 788–798. PMLR (2020)