

MULTI-FIDELITY STOCHASTIC TRUST REGION METHOD WITH ADAPTIVE SAMPLING

Yunsoo Ha¹ and Juliane Mueller¹

¹Computational Science Center, National Renewable Energy Laboratory, 15013 Denver West Parkway, Golden, 80401, Colorado, USA

ABSTRACT

Simulation optimization is often hindered by the high cost of running simulations. Multi-fidelity methods offer a promising solution by incorporating cheaper, lower-fidelity simulations to reduce computational time. However, the bias in low-fidelity models can mislead the search, potentially steering solutions away from the high-fidelity optimum. To overcome this, we propose ASTRO-MFDF, an adaptive sampling trust-region method for multi-fidelity simulation optimization. ASTRO-MFDF features two key strategies: (i) it adaptively determines the sample size and selects appropriate sampling strategies to reduce computational cost; and (ii) it selectively uses low-fidelity information only when a high correlation with the high-fidelity is anticipated, reducing the risk of bias. We validate the performance and computational efficiency of ASTRO-MFDF through numerical experiments using the SimOpt library.

1 INTRODUCTION

Simulation optimization (SO) has become a key method for optimizing objective functions in uncertain environments, gaining significant attention for its ability to tackle real-world problems involving randomness and complex systems, such as quantum computing and renewable energy (Ha et al. 2025; Sakki et al. 2022). However, its practical implementation can be challenging due to high computational cost of evaluating a stochastic function value. To address this, simulation models can be developed at different levels of fidelity, which sometimes requires substantial modeling or coding effort, a method known as multi-fidelity (MF) simulation. In MF simulation, models are designed with a hierarchical structure, where high-fidelity models provide detailed and accurate representations of the process, while low-fidelity models offer a more computationally efficient, simplified version. For example, a high-fidelity model may represent the full manufacturing process in detail, whereas lower-fidelity models might omit certain machines that are not critical (Zhang et al. 2022). Another way to construct MF simulation models is by modifying the length of the simulation runs. When aiming to optimize systems under a steady-state condition, using shorter run lengths results in lower-fidelity models that produce less accurate output estimates but require reduced computational effort (Chen et al. 2017).

In this paper, we consider the multi-fidelity simulation optimization (MFSO) problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} f^0(\mathbf{x}) := \mathbb{E}_{\Xi^0}[F^0(\mathbf{x}, \xi^0)], \quad (1)$$

where $f^0 : \mathbb{R}^d \rightarrow \mathbb{R}$ is nonconvex and has a lower bound, $F^0 : \mathbb{R}^d \times \Xi^0 \rightarrow \mathbb{R}$ is a random function, and $\xi^0 : \Omega \rightarrow \Xi^0$ is a random element. Here, the index 0 represents the highest-fidelity simulation, while increasing index values correspond to lower-fidelity simulations. We consider zeroth-order stochastic oracles, where the derivative information is not directly available from the Monte Carlo simulation. We allow each realization $F^0(\mathbf{x}, \xi_i^0)$ to be nonconvex and nonsmooth, as is common in complex simulation models. Since we only have access to $F^0(\mathbf{x})$, we can estimate $f^0(\mathbf{x})$ by $\bar{F}^0(\mathbf{x}, n) = n^{-1} \sum_{i=1}^n F^0(\mathbf{x}, \xi_i^0)$, variance of $F^0(\mathbf{x}, \xi^0)$ by $(\hat{\sigma}^0(\mathbf{x}, n))^2 = n^{-1} \sum_{i=1}^n (F^0(\mathbf{x}, \xi_i^0) - \bar{F}^0(\mathbf{x}, n))^2$, and covariance between $F^0(\mathbf{x}, \xi^0)$ and $F^i(\mathbf{x}, \xi^i)$ by $\hat{\sigma}^{0,i}(\mathbf{x}, n) = n^{-1} \sum_{j=1}^n (F^0(\mathbf{x}, \xi_j^0) - \bar{F}^0(\mathbf{x}, n))(F^i(\mathbf{x}, \xi_j^i) - \bar{F}^i(\mathbf{x}, n))$ for any $i \in \mathbb{N}$. Typically,

iterative algorithms are used to solve (1) by generating a random sequence of iterates $\{\mathbf{X}_k\}$ that converges toward an optimal solution. In the context of MFSO, the goal is to accelerate this convergence by effectively leveraging information from MF simulations.

In deterministic MF problems, multi-fidelity Bayesian optimization (MFBO) is widely used, with co-kriging serving as the surrogate model (Do and Zhang 2024). However, applying MFBO to solve (1) requires two key assumptions (Foumani et al. 2023). First, function estimates must be sufficiently accurate. Second, lower-fidelity functions should exhibit a strong correlation with the high-fidelity function throughout the entire search space. Otherwise, lower-fidelity data can negatively impact the co-kriging model, slowing down optimization even more than using only the high-fidelity model. To mitigate these challenges, it is crucial to carefully select both the sample size (n) and the regions where lower-fidelity data provides meaningful information. This consideration leads to the following research question:

(RQ) *Where, which fidelity, and how many times should oracles be queried to efficiently solve (1)?*

The adaptive sampling trust region method (ASTRO-DF) is one of the most effective algorithms for addressing this question, as it dynamically adjusts the sample size and trust region (Shashaani et al. 2018). Specifically, the sample size is chosen to balance the trade-off between the optimality gap and estimation errors, helping to reduce the computational burden. In addition, the search space is confined to a neighborhood of the current iterate, known as the trust region, by constructing and optimizing a surrogate model within this region. The trust region is updated based on whether the surrogate model provides a sufficiently accurate approximation of the true objective function, enabling correlation measurement of the MF functions within a localized search space. See (Conn et al. 2009) for the details of deterministic trust region method for derivative-free optimization. To address (RQ) in the bi-fidelity setting, we recently proposed an extension of ASTRO-DF, known as ASTRO-BFDF (Ha and Mueller 2024). ASTRO-BFDF differs from ASTRO-DF in two primary aspects: i) it maintains two separate trust regions for high-fidelity and low-fidelity functions, and ii) it employs an adaptive sampling approach that utilizes both Bi-fidelity Monte Carlo (BFMC) and standard Monte Carlo (MC) to reduce the computational burden.

BFMC is a special case of the multi-fidelity Monte Carlo (MFMC) method that uses exactly two levels of simulation fidelity. MFMC has been introduced in (Peherstorfer et al. 2016) to enhance the efficiency of Monte Carlo methods by leveraging MF simulation oracles through a control variate approach. Specifically, the MFMC estimate is obtained by

$$\hat{F}^t(\mathbf{x}, \mathbf{n}, \mathbf{c}) = \frac{1}{n^t} \sum_{j=1}^{n^t} F^t(\mathbf{x}, \xi_j^t) + \sum_{i=t+1}^q c^i \left(\frac{1}{n^i} \sum_{j=1}^{n^i} F^i(\mathbf{x}, \xi_j^i) - \frac{1}{n^{i-1}} \sum_{j=1}^{n^{i-1}} F^i(\mathbf{x}, \xi_j^i) \right), \quad (2)$$

where q denotes the index of the lowest-fidelity simulation, t is the target fidelity level, $\mathbf{c} = \{c^1, c^2, \dots, c^q\}$ with c^i indicating the coefficient associated with i -th fidelity oracle, and $\mathbf{n} = \{n^0, n^1, \dots, n^q\}$ with n^i denoting the sample size used for the i -th fidelity oracle. Given fixed \mathbf{n} and \mathbf{c} such that $n^0 < n^1 < \dots < n^q$, and assuming Common Random Numbers (CRN) coupling across fidelity levels, the MFMC estimate is unbiased estimate for $f^0(\mathbf{x})$ with a variance

$$\text{Var}(\hat{F}^t(\mathbf{x}, \mathbf{n}, \mathbf{c})) = \frac{(\sigma^t(\mathbf{x}))^2}{n} + \sum_{i=t+1}^q \left(\frac{1}{n^{i-1}} - \frac{1}{n^i} \right) ((c^i)^2 (\sigma^i(\mathbf{x}))^2 - 2c^i \sigma^{0,i}(\mathbf{x})), \quad (3)$$

where $(\sigma^i(\mathbf{x}))^2$ is the variance of $F^i(\mathbf{x}, \xi^i)$ and $\sigma^{0,i}$ is the covariance between $F^0(\mathbf{x}, \xi^0)$ and $F^i(\mathbf{x}, \xi^i)$. Suppose that we set $t = 0$ and $q = 1$. Then the variance reduction is achieved when the condition $2c^1 \sigma^{0,1}(\mathbf{x}) \geq (c^1)^2 (\sigma^1(\mathbf{x}))^2$ holds. The CRN assumption typically induces positive correlation between fidelity levels, making this condition feasible; later we propose a principled choice of \mathbf{c} to minimize the variance of the MFMC estimator.

In this paper, we extend ASTRO-BFDF to address the MFSO problem (1) by incorporating the MFMC technique, which we refer to as ASTRO-MFDF. We first present a multi-fidelity adaptive sampling algorithm

in Section 2. Then, we introduce the ASTRO-MFDF algorithm in Section 3. Finally, we demonstrate the performance of ASTRO-MFDF through experiments with stochastic Rosenbrock functions and continuous (s, S) inventory problems in Section 4. Throughout the paper, we use capital letters for random variables and bold font for vectors.

2 MULTI-FIDELITY ADAPTIVE SAMPLING

The adaptive sampling method used within ASTRO-DF performs sequential estimation to effectively reduce computational effort associated with MC. Specifically, the sample size at $\mathbf{x} \in \mathbf{R}^d$ has been determined by

$$N_k(\mathbf{x}) = \min \left\{ n \in \mathbb{N} : \underbrace{\frac{\max\{\sigma_{lb}, \hat{\sigma}^0(\mathbf{x}, n)\}}{\sqrt{n}}}_{\text{stochastic error}} \leq \underbrace{\frac{\kappa \Delta_k^2}{\sqrt{\lambda_k}}}_{\text{optimality gap}} \right\}, \quad (4)$$

where κ , λ_k , and σ_{lb} are positive constants, Δ_k is the trust region size at iteration k (Ha, Shashaani, and Pasupathy 2024). The condition in (4) balances the stochastic error and the optimality gap, represented by the standard deviation estimate and the square of the trust region size respectively. Following the same principle, when employing the MFMC estimate, our goal is to determine \mathbf{n} and \mathbf{c} such that the condition in (4) is satisfied by replacing the stochastic error term under MC, i.e., $n^{-1}(\hat{\sigma}^0(\mathbf{x}, \mathbf{n}))^2$, with that under MFMC, i.e., the estimate of $\text{Var}(\hat{F}^0(\mathbf{x}, \mathbf{n}, \mathbf{c}))$. However, incorporating the adaptive sampling strategy to MFMC introduces the following two challenges:

- (C1) The decision variables become multi-dimensional, consisting of the vectors \mathbf{n} and \mathbf{c} . In the MC case, the decision variable is one-dimensional, specifically the scalar n . Thus, to determine $N_k(\mathbf{x})$, we initially compute $\hat{\sigma}^0(\mathbf{x}, n=3)$ and increment n by one until the condition in (4) is satisfied, due to the unknown variance $(\sigma^0(\mathbf{x}))^2$. Similarly, the variances and covariances appearing in (3) are also unknown, necessitating sequential estimation for these quantities as well. However, designing a straightforward sequential estimation algorithm to determine \mathbf{n} and \mathbf{c} while minimizing computational burden is challenging.
- (C2) Considering factors such as the variances and covariances presented in (3) and the querying costs of the MF simulation oracles, the computational cost of MFMC can be higher than that of MC for achieving the same level of accuracy, despite MFMC being a variance reduction method. Hence, we must determine which MC method (MC or MFMC) to use while sequentially estimating variances and covariances, introducing an additional binary decision variable.

In (C1), we should determine the optimal \mathbf{n} and \mathbf{c} based on the available information. Suppose, for example, that variance and covariance estimates are currently available with sample sizes \tilde{n}^i for the i -th fidelity simulation for any $\tilde{n}^i \geq 2$. Using these estimates, we can obtain an approximation of (3), specifically, $\widehat{\text{Var}}(\hat{F}^0(\mathbf{x}, \mathbf{n}, \mathbf{c}))$, by replacing $\sigma^i(\mathbf{x})$ and $\sigma^{0,i}(\mathbf{x})$ with their corresponding estimators $\hat{\sigma}^i(\mathbf{x}, \tilde{n}^i)$ and $\hat{\sigma}^{0,i}(\mathbf{x}, \tilde{n}^0)$, respectively. We can then determine the optimal solutions for \mathbf{n} and \mathbf{c} , denoted by \mathbf{n}_* and \mathbf{c}_* , by solving the following optimization problem:

$$\begin{aligned} [\mathbf{n}_*, \mathbf{c}_*] \in \arg \min_{\mathbf{n}, \mathbf{c} \in \mathbf{R}^q} & \sum_{i=0}^q w^i n^i \\ \text{subject to} & \widehat{\text{Var}}(\hat{F}^0(\mathbf{x}, \mathbf{n}, \mathbf{c})) \leq \kappa^2 \Delta_k^4 \lambda_k^{-1} \\ & n^i - n^{i+1} \leq 0 \quad \forall i \in \{0, 1, \dots, q-1\} \\ & \tilde{n}^i - n^i \leq 0 \quad \forall i \in \{0, 1, \dots, q\}, \end{aligned} \quad (5)$$

where $\mathbf{w} = \{w^0, w^1, \dots, w^q\}$ with w^i indicating the cost of querying the i -th fidelity simulation oracle. The first constraint originates from the condition stated in (4). The predicted optimal cost of the MFMC method

then becomes $\sum_{i=0}^q w^i n_*^i$. Additionally, the predicted optimal cost for the MC method is calculated as $w^0 n_p$, where $n_p = \lceil (\hat{\sigma}^0(\mathbf{x}, \tilde{\mathbf{n}}^0))^2 \lambda_k (\kappa^2 \Delta_k^4)^{-1} \rceil$, based on the condition from (4). Therefore, MFMC is selected if $\sum_{i=0}^q w^i n_*^i \leq w^0 n_p$; otherwise, MC is chosen. This resolves the issue mentioned in (C2). It is worth noting that problem (5) is a nonlinear continuous optimization problem with mostly linear constraints and a single nonlinear constraint, where the problem dimension is relatively small. Since the primary computational cost in simulation optimization lies in the simulation itself, the cost of solving problem (5) is assumed to be negligible. We also note that the optimal integer sample sizes can be obtained by rounding up n_* .

If the chosen estimate is sufficiently accurate—specifically, if it satisfies the condition in (4) based on its own variance estimate—then that function estimate can directly serve as the final output of the multi-fidelity adaptive sampling method. If this accuracy condition is not met, further replications are required to refine the estimate. In this case, when MC has been chosen, the 0-th fidelity oracle naturally becomes the default choice. Conversely, if MFMC is more economical but the accuracy condition $\widehat{\text{Var}}(\hat{F}^0(\mathbf{x}, \tilde{\mathbf{n}}, \mathbf{c}_*)) > \kappa^2 \Delta_k^4 \lambda_k^{-1}$ remains unmet, i.e., $\tilde{n}^j < n_*^j - 1$ for some $j \in \{0, 1, \dots, q\}$, we proceed by querying the highest such fidelity level, as allocating additional samples to higher fidelity levels is generally more effective in reducing the variance. After selecting the appropriate fidelity level, we conduct additional replications, update variance and covariance estimates accordingly, and then resolve the optimization problem (5). This iterative process continues until one of the accuracy conditions is ultimately satisfied. Although we estimated f^0 above, target fidelity can be any $i \in \{0, 1, \dots, q\}$.

To complement the algorithmic description, we include a flowchart that illustrates the overall logic of the multi-fidelity adaptive sampling method (MFAS), including the decision-making process between MC and MFMC, as well as the adaptive sample size update based on accuracy checks (see Figure 1).

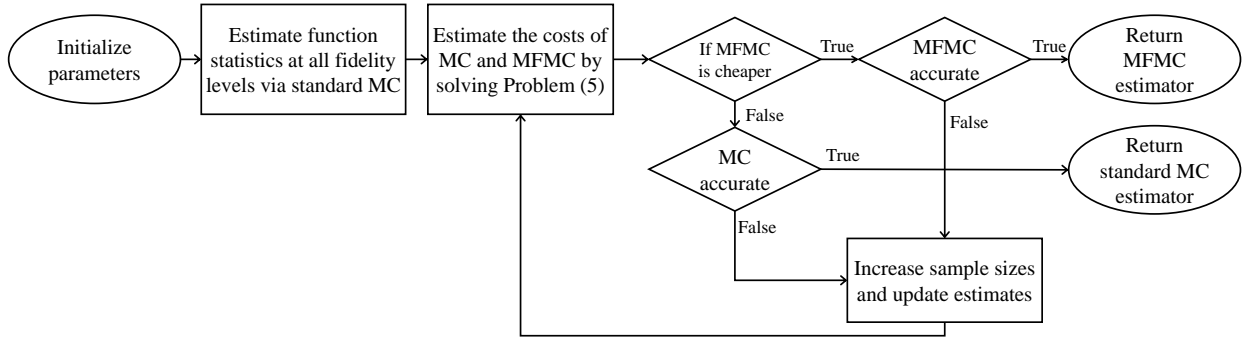


Figure 1: Flowchart of MFAS. The algorithm first initializes relevant parameters and estimates function statistics (mean, variance, and covariance) at all fidelity levels using standard MC sampling. The algorithm then solves Problem (5) to assess the cost-effectiveness of MFMC versus standard MC estimators. If the cheaper estimator also satisfies the required accuracy, it is selected and returned. Otherwise, the sample sizes are increased, and the estimates are updated iteratively.

3 MULTI-FIDELITY STOCHASTIC TRUST REGION METHOD WITH MFAS

In traditional stochastic trust region methods (Chang, Hong, and Wan 2013; Chen, Menickelly, and Scheinberg 2018), a single interpolation/regression model (M_k) is constructed at iteration k , and the next candidate point is determined by approximately minimizing this model within the trust region, usually by computing a Cauchy point or using iterative methods such as the conjugate gradient method; see (Nocedal and Wright 2006) for further details. The key distinction of ASTRO-MFDF is that, instead of relying on a single model, it can construct multiple interpolation models, denoted as M_k^i for any $i \in \{0, 1, \dots, q\}$, across multiple trust regions, represented by $\Delta_k := \{\Delta_k^0, \Delta_k^1, \dots, \Delta_k^q\}$. The local model constructed with lower-fidelity oracles is prioritized in selecting the next iterate, effectively guiding the iterates closer to the optimal

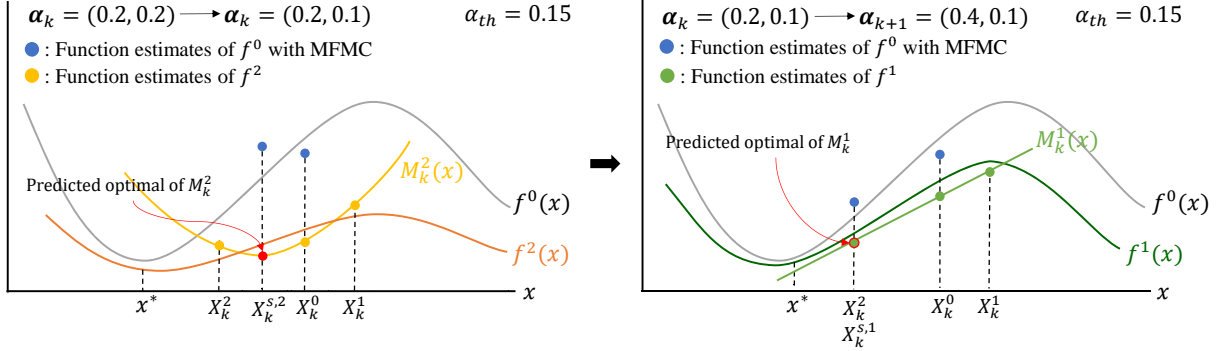


Figure 2: Illustration of the inner loop (Steps 2–9) of ASTRO-MFDF, which iteratively applies Algorithm 2 across fidelity levels. Since M_k^2 fails to generate a better candidate than X_k^0 , α_k^2 decreases, and M_k^1 is instead utilized, successfully yielding a better candidate. As a result, at iteration $k+1$, M_k^1 will be constructed first based on the updated α_{k+1} .

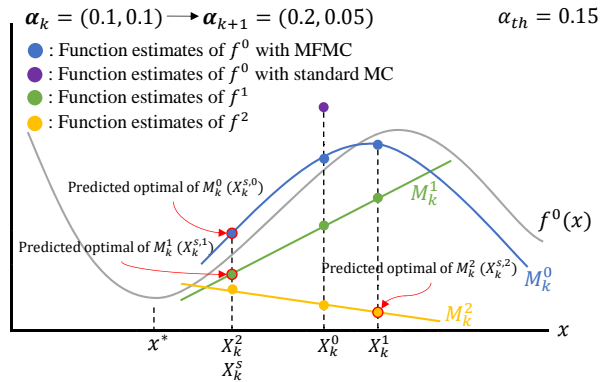


Figure 3: Illustration of Steps 10–15 in ASTRO-MFDF. Since all $\alpha_k^i < \alpha_{th}$, design points are shared across all models. Lower-fidelity estimates (green and yellow dots) are used in the MFMC estimator (blue dots) to approximate f^0 and are then reused in constructing M_k^1 and M_k^2 . As M_k^2 fails to improve over the incumbent, α_k^2 decreases, while M_k^1 succeeds, increasing α_k^1 .

solution using only lower-fidelity simulations. However, constructing q local models at every iteration can be cumbersome and a waste of computational resources, especially since lower-fidelity simulations may not yield better solutions in certain feasible regions due to the inherent bias between the high-fidelity function and the lower-fidelity functions. Hence, we employ an adaptive correlation vector $\alpha_k := \{\alpha_k^1, \dots, \alpha_k^q\}$, updated dynamically to capture the correlation between f^0 and f^i for any $i \in \{1, 2, \dots, q\}$. Specifically, if the candidate point generated by M_k^i achieves a sufficient reduction in the estimated objective function, α_k^i increases; otherwise, it decreases. When α_k^i exceeds a threshold α_{th} (See Figure 2), M_k^i is constructed using the individual design set, as a strong correlation between f^0 and f^i is expected. If $\alpha_k^i < \alpha_{th}$ for all $i \in \{1, 2, \dots, q\}$ (See Figure 3), the main local model defaults to M_k^0 for finding the candidate point. Meanwhile, M_k^i is constructed using same design points for 0-th fidelity simulation with the main purpose of updating α_k^i , even though some design points may lie outside $\mathcal{B}(\mathbf{X}_k, \Delta_k^i)$, the trust region centered at \mathbf{X}_k with radius Δ_k^i . It is important to note that constructing M_k^i with same design points requires minimal additional computational effort, as numerous replications of the i -th fidelity simulation are already available from the

Algorithm 1 ASTRO-MFDF

Input: Initial incumbent $\mathbf{x}_0 \in \mathbb{R}^d$, initial and maximum trust region radius Δ_0 and $\Delta_{\max} > 0$, model fitness thresholds $0 < \eta < 1$ and certification threshold $\mu > 0$, expansion and shrinkage constants $\gamma_1 > 1$ and $\gamma_2 \in (0, 1)$, sample size lower bound sequence $\{\lambda_k\} = \{\mathcal{O}(\log k)\}$, adaptive sampling constant $\kappa > 0$, correlation vector $\boldsymbol{\alpha}_0 = \{\alpha_0^1, \alpha_0^2, \dots, \alpha_0^q\}$, and lower bound of an initial variance approximation $\sigma_{lb} > 0$, and sufficient reduction constant $\zeta > 0$.

- 1: **for** $k = 0, 1, 2, \dots$ **do**
 - 2: **for** $t = q, q - 1, \dots, 1$ **do**
 - 3: Obtain \mathbf{X}_k^s and Δ_k^t by calling Algorithm 2 (ASTRO-LFDF- t).
 - 4: **if** Algorithm 2 yielded a candidate with sufficient HF function reduction (i.e., $\mathbf{X}_k^s \neq \mathbf{X}_k^0$) **then**
 - 5: Set $(\mathbf{X}_{k+1}, \Delta_{k+1}^t) = (\mathbf{X}_k^s, \gamma_1 \Delta_k^t)$ and $\alpha_{k+1}^t = \gamma_1 \alpha_k^t$.
 - 6: Set $\Delta_{k+1}^j = \max\{\Delta_k^j, \Delta_{k+1}^t\}$ for all $j \in \{0, 1, \dots, t - 1\}$ and $k = k + 1$.
 - 7: **break**
 - 8: **end if**
 - 9: **end for**
 - 10: **if** Algorithm 2 failed to find a better candidate than the current solution (i.e., $\mathbf{X}_k^s = \mathbf{X}_k^0$) **then**
 - 11: Select $\{\mathbf{X}_k^i\}_{i=0}^{2d} \subset \mathcal{B}(\mathbf{X}_k; \Delta_k^0)$.
 - 12: Estimate the t -fidelity function \tilde{F}^t at $\{\mathbf{X}_k^i\}_{i=0}^{2d}$ using MFAS with $\Delta_k = \Delta_k^0$, construct local models $M_k^t(\mathbf{X})$, approximately compute the minimizers $\mathbf{X}_k^{s,t} \in \arg \min_{\|\mathbf{X} - \mathbf{X}_k\| \leq \Delta_k^t} M_k^t(\mathbf{X})$ and then, estimate the 0-fidelity function $\tilde{F}^0(\mathbf{X}_k^{s,t})$ using MFAS with $\Delta_k = \Delta_k^0$ for all $t \in \{0, 1, \dots, q\}$.
 - 13: Set the candidate point $\mathbf{X}_k^s \in \arg \min_{\mathbf{x} \in \{\mathbf{X}_k^{s,i}\}_{i=0}^q} \tilde{F}^0(\mathbf{x})$, and compute the success ratio $\hat{\rho}_k$ and $\hat{\rho}_k^t$ for any $t \in \{1, 2, \dots, q\}$ as

$$\hat{\rho}_k = \frac{\tilde{F}^0(\mathbf{X}_k^0) - \tilde{F}^0(\mathbf{X}_k^s)}{M_k^0(\mathbf{X}_k^0) - M_k^0(\mathbf{X}_k^{s,0})} \text{ and } \hat{\rho}_k^t = \frac{\tilde{F}^0(\mathbf{X}_k^0) - \tilde{F}^0(\mathbf{X}_k^{s,t})}{\max\{\zeta(\Delta_k^0)^2, M_k^0(\mathbf{X}_k^0) - M_k^0(\mathbf{X}_k^{s,t})\}}.$$
 - 14: **If** M_k^t succeeds (i.e., $\hat{\rho}_k^t \geq \eta$), set $\alpha_{k+1}^t = \gamma_1 \alpha_k^t$; otherwise set $\alpha_{k+1}^t = \gamma_2 \alpha_k^t$ for all $t \in \{1, 2, \dots, q\}$.
 - 15: Set $(\mathbf{X}_{k+1}, \Delta_{k+1}^0) = \begin{cases} (\mathbf{X}_k^s, \min\{\gamma_1 \Delta_k^0, \Delta_{\max}\}) & \text{if } \hat{\rho}_k \geq \eta \text{ and } \mu \|\nabla M_k^0(\mathbf{X}_k)\| \geq \Delta_k^0, \\ (\mathbf{X}_k, \gamma_2 \Delta_k^0) & \text{otherwise.} \end{cases}$
 - 16: Set $\Delta_{k+1}^t = \min\{\Delta_k^t, \Delta_k^0\}$ for all $t \in \{1, 2, \dots, q\}$, and $k = k + 1$.
 - 17: **end if**
 - 18: **end for**
-

MFMC estimates in MFAS during the construction of M_k^0 . To provide an overview of the ASTRO-MFDF procedure, we include a flowchart summarizing the main logic, including the role of the ASTRO-LFDF subroutine, the adaptive update of $\boldsymbol{\alpha}_k$, and the candidate selection (Figure 4). While the flowchart depicts the bi-fidelity setting as a special case, the logic extends naturally to the general multi-fidelity setting. The pseudocode of ASTRO-MFDF is listed in Algorithm 1.

We now outline some minor details of ASTRO-MFDF for practical purposes. First, we ensure that $\Delta_k^i \geq \Delta_k^j$ for any $i < j$. In practice, the main computational challenge arises from the increasing sample size, which scales at a rate of Δ_k^{-4} as Δ_k converges to zero. Therefore, maintaining a large trust region for higher-fidelity functions while progressing toward better iterates with a relatively smaller trust region for lower-fidelity functions enhances finite-time efficiency. Second, when low correlation between f^0 and f^i for all $i \in \{1, 2, \dots, q\}$ is predicted, i.e., $\alpha_k^i < \alpha_{th}$ for all $i \in \{1, 2, \dots, q\}$, we generate $q + 1$ candidate points $(\mathbf{X}_k^{s,q})$ by minimizing M_k^i within Δ_k^0 for each $i \in \{0, 1, \dots, q\}$ (Step 12 in Algorithm 1). The final

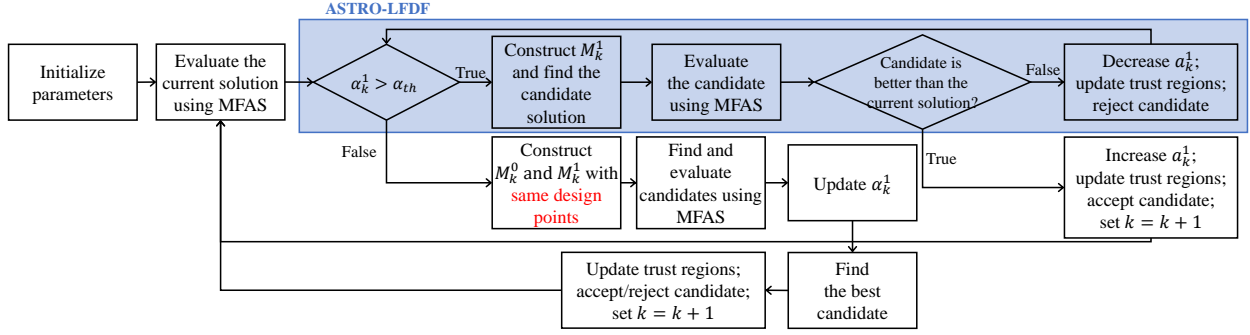


Figure 4: Flowchart of ASTRO-MFDF, illustrating the integration of ASTRO-LFDF (highlighted in blue), the decision process based on the correlation measure, the generation of candidate solutions, and the update of trust regions and correlation values. For clarity, the flowchart depicts the bi-fidelity case ($q = 1$) as a special instance.

candidate (\mathbf{X}_k^s) is then selected as the point with the lowest function estimate, maximizing the likelihood of a successful iteration. Lastly, for an iteration with a candidate point from M_k^t for any $t \in \{1, 2, \dots, q\}$ to be considered successful, we impose the additional condition $\tilde{F}^0(\mathbf{X}_k^0) - \tilde{F}^0(\mathbf{X}_k^{s,q}) \geq \zeta \eta (\Delta_k^0)^2$ to prevent success arising from a negligible model reduction $M_k^t(\mathbf{X}_k^0) - M_k^t(\mathbf{X}_k^{s,t})$ (Step 9 in Algorithm 2).

Algorithm 2 $[\mathbf{X}_k^s, \Delta_k^t] = \text{ASTRO-LFDF-}t$

Input: $\mathbf{X}_k^0, \Delta_k^t$, model fitness thresholds $0 < \eta < 1$, expansion and shrinkage constants $\gamma_1 > 1$ and $\gamma_2 \in (0, 1)$, sample size lower bound sequence $\{\lambda_k\} = \{\mathcal{O}(\log k)\}$, adaptive sampling constant $\kappa > 0$, correlation constant $\alpha_k^t > 0$, correlation threshold $\alpha_{th} > 0$, lower bound of an initial variance approximation $\sigma_{lb} > 0$, and sufficient reduction constant $\zeta > 0$.

- 1: **loop**
 - 2: **if** $\alpha_k^t < \alpha_{th}$ **then**
 - 3: Set $\mathbf{X}_k^{s,t} = \mathbf{X}_k^0$
 - 4: **break**
 - 5: **end if**
 - 6: Select the design set $\{\mathbf{X}_k^i\}_{i=0}^{2d} \subset \mathcal{B}(\mathbf{X}_k; \Delta_k^t)$.
 - 7: Estimate t -fidelity function $\tilde{F}^t(\mathbf{X}_k^i)$ at $\{\mathbf{X}_k^i\}_{i=0}^{2d}$ using MFAS with $\Delta_k = \Delta_k^t$, construct local model $M_k^t(\mathbf{X})$, and approximately compute $\mathbf{X}_k^{s,t} = \arg \min_{\|\mathbf{X} - \mathbf{X}_k\| \leq \Delta_k^t} M_k^t(\mathbf{X})$.
 - 8: Estimate the 0-fidelity function $\tilde{F}_k^0(\mathbf{X}_k^{s,t})$ and $\tilde{F}_k^0(\mathbf{X}_k^0)$ using MFAS with $\Delta_k = \Delta_k^t$.
 - 9: Compute the success ratio $\hat{\rho}_k = \tilde{F}_k^0(\mathbf{X}_k^0) - \tilde{F}_k^0(\mathbf{X}_k^{s,t}) / \max\{\zeta (\Delta_k^0)^2, M_k^t(\mathbf{X}_k^0) - M_k^t(\mathbf{X}_k^{s,t})\}$.
 - 10: **if** $\hat{\rho}_k \geq \eta$ **then**
 - 11: **break**
 - 12: **end if**
 - 13: Set $\Delta_k^t = \gamma_2 \Delta_k^t$ and $\alpha_k^t = \gamma_2 \alpha_k^t$
 - 14: **end loop**
 - 15: **return** $[\mathbf{X}_k^s = \mathbf{X}_k^{s,t}, \Delta_k^t]$
-

4 NUMERICAL EXPERIMENTS

In this section, we analyze the finite-time performance of solvers such as ASTRO-MFDF, ASTRO-DF, and Nelder-Mead, across two different problems. First, we test the solvers on a stochastic Rosenbrock function.

Second, we use continuous (s, S) inventory problems to evaluate performance in steady-state simulation optimization, where s is a threshold for ordering and S is a maximum inventory level.

The experiments in this study were carried out using SimOpt, a benchmarking framework designed for evaluating simulation optimization algorithms (Eckman et al. 2023). The experimental process consists of two stages. In the first stage, we generate 20 independent optimization runs (macro-replications) for each solver on each problem, where each run represents a distinct stochastic trajectory. Every solver is allocated a fixed budget specific to the problem at hand and must strategically decide how to utilize this budget, including selecting locations, fidelity levels, and sample sizes. For instance, if the budget is set to a specific number, such as B , it allows for B queries to the simulation oracle to solve the given problem. In the MF scenario, we introduce a cost ratio vector (\mathbf{w}) , whose component satisfies $w^0 = 1$ for the highest fidelity and decrease thereafter, i.e., $w^0 > w^1 > \dots$. Therefore, if we query 10 oracles at fidelity level 0 and 20 oracles at fidelity level 1, the total budget consumed is calculated as $10w^0 + 20w^1$. Throughout each run, the solver estimates objective values at various solutions (design points), using a chosen number of replications to guide its search. The second stage involves evaluating the recommended solutions from each run using 200 additional simulations to estimate the true performance of the solutions.

4.1 Experiments on stochastic Rosenbrock function

The deterministic MF Rosenbrock functions were introduced in (Mainini et al. 2022) as $f^0(\mathbf{x}) = \sum_{i=1}^d 10(x_{i+1} - x_i^2)^2 + (1 - x_i)^2$, $f^1(\mathbf{x}) = \sum_{i=1}^{d-1} 5(x_{i+1} - x_i^2)^2 + (-2 - x_i)^2 - \sum_{i=1}^d 0.5x_i$, and $f^2(\mathbf{x}) = (f^0 - 4 - \sum_{i=1}^d 0.5x_i)(10 + \sum_{i=1}^d 0.25x_i)^{-1}$. See Figure 5 for two-dimensional loss landscapes. To introduce stochasticity, noise terms E_i^t (drawn from a normal distribution with zero mean and dimension-dependent variance for any $i \in \{1, 2, \dots, d\}$ and $t \in \{0, 1, 2\}$) are added: $\sum_{i=1}^d E_i^0$ to f^0 , $\sum_{i=1}^d (E_i^0 + E_i^1)/2$ to f^1 , and $\sum_{i=1}^d (E_i^0 + E_i^2)/2$ to f^2 . The cost of querying each fidelity oracle is given by $\mathbf{w} = (1, 0.3, 0.1)$.




Figure 5: Contour map of the deterministic Rosenbrock MF function for $d = 2$. The optimal solution is $(1, 1)$ for f^0 , while the approximate optima for the lower-fidelity functions are $(-1.4, 2)$ for f^1 and $(1.4, 2)$ for f^2 with a box constraint $-2 \leq x_i \leq 2$ for all $i \in \{1, 2\}$.

To illustrate the mechanism of ASTRO-MFDF, we compare its optimization trajectory with that of ASTRO-DF on the 2-dimensional stochastic Rosenbrock function (Figure 6). As shown in Figure 6b, the optimization with ASTRO-MFDF begins using the lowest-fidelity function f^2 , indicated in red. After several iterations, the local models constructed with lower-fidelity functions (M_k^1 and M_k^2) fail to yield improved iterates, prompting the algorithm to rely more heavily on the highest-fidelity function, represented in black. Later in the process, M_k^2 starts to produce better solutions again (red) enabling the algorithm to reduce computational cost per iteration. As a result, ASTRO-MFDF completes 24 iterations within the 500 0-th fidelity budget, compared to just 11 iterations by ASTRO-DF. It is worth noting that M_k^1 has not

been used, as the correlation between f^0 and f^1 is expected to be low—i.e., α_k^1 remains small—along the sample path of $\{\mathbf{X}_k\}$ generated by ASTRO-MFDF (see the contour map of f^0 and f^1 in Figure 5). This illustrates how ASTRO-MFDF can dynamically select the appropriate fidelity level at each iterate to maximize the computational efficiency of the MF approach. We also tested higher-dimensional problems, as shown in Figure 7. As the dimensionality increases, ASTRO-MFDF demonstrates faster convergence, since constructing local models using lower-fidelity functions becomes significantly more cost-effective than using high-fidelity functions in high-dimensional settings.

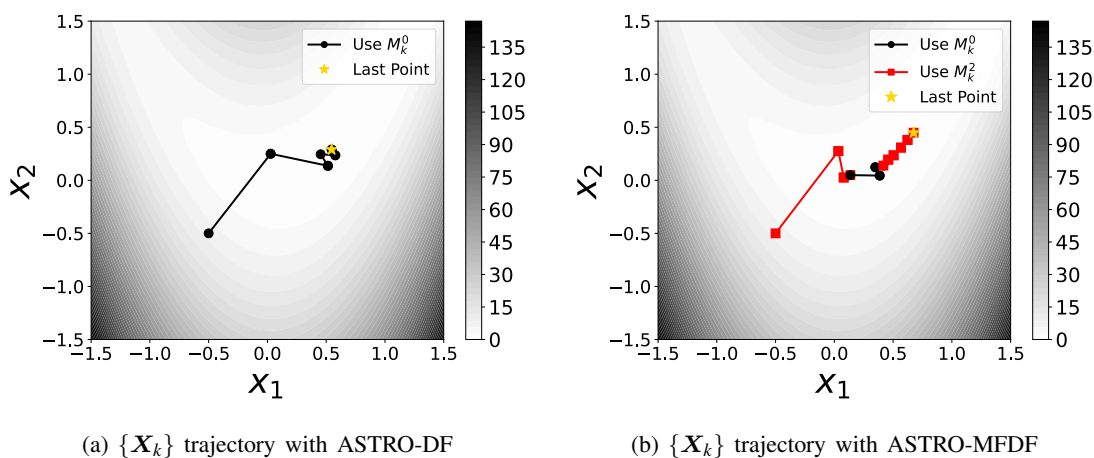


Figure 6: One sample path of $\{\mathbf{X}_k\}$ with ASTRO-DF and ASTRO-MFDF on the 2-dimensional stochastic Rosenbrock function. Starting from the initial point $(-0.5, -0.5)$ with a budget of 500 0-th fidelity oracle calls, the sequence $\{\mathbf{X}_k\}$ converges to $(0.55, 0.29)$ in (a), with a corresponding f^0 value of approximately 0.308, and to $(0.67, 0.45)$ in (b), with a corresponding f^0 value of approximately 0.109.




Figure 7: Finite-time performance on the stochastic Rosenbrock function, with a 95% confidence interval with initial design point $(-0.5) \times d$. The x-axis represents a budget of 5000 0-th fidelity oracle calls, and the y-axis shows the objective function value on a logarithmic scale.

4.2 Experiments on continuous inventory problem

In this section, we consider the (s, S) inventory problem with continuous decision variables to create a more realistic experimental setting. The objective is to determine the optimal values of s and S that minimize the expected total cost, which includes holding costs, ordering costs, and backorder costs. Uncertainty in the system arises from two sources: (1) demand in each period follows an exponential distribution with mean θ , and (2) lead times are drawn from a Poisson distribution with mean ℓ periods. For further details, see (Eckman, Henderson, Shashaani, and Pasupathy 2023). In the MF simulation, the 0-fidelity model runs for 100 days, the 1-fidelity model for 50 days, and the 2-fidelity model for 30 days, which implies the cost ratio vector $\mathbf{w} = (1, 0.5, 0.3)$.




Figure 8: Contour map of the estimated expected total cost for the (s, S) inventory problem with $\theta = 400$ and $\ell = 3$. Since the true objective function is unknown, the plots are based on estimates obtained from 10 samples.

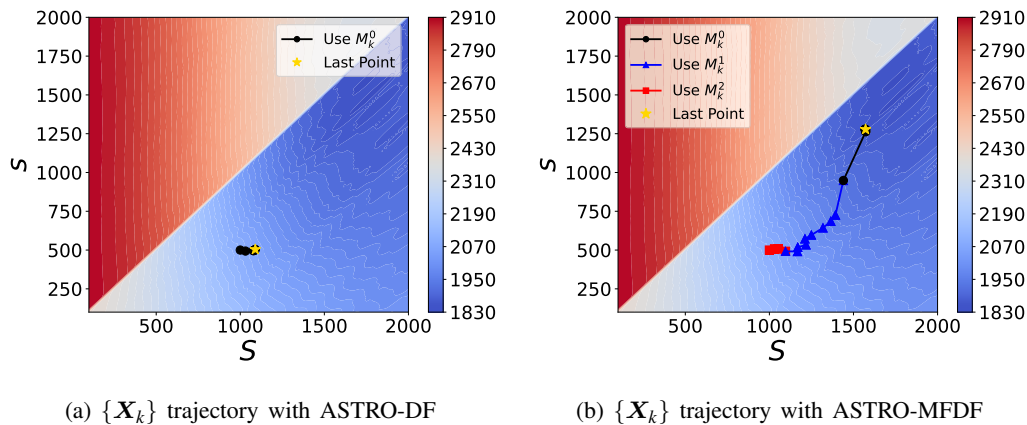


Figure 9: One sample path of $\{X_k\}$ with ASTRO-DF and ASTRO-MFDF on the continuous (s, S) inventory problem with $\theta = 400$ and $\ell = 3$. Starting from the initial point $(500, 1000)$ with a budget of 1000 oracle calls, the sequence $\{X_k\}$ converges to $(504.13, 1089.27)$ in (a), and to $(1277.85, 1571.87)$ in (b). The contour map shows the estimated expected total cost with 100 samples.

As shown in Figure 8, the estimated objective function with a small sample size appears highly non-convex and non-smooth due to stochastic noise. As a result, the iterates are more likely to get stuck in regions that may not even correspond to local minima of the true objective function. When this happens, it typically indicates that the step size or search space has already become too small, making it difficult

to escape without a significant amount of additional computational effort. As it can be seen in Figure 9, ASTRO-MFDF naturally addresses this issue by utilizing lower-fidelity models to escape suboptimal regions and preserving a large trust region for the high-fidelity function, thereby minimizing unnecessary computational costs. We also explored a range of settings for θ and ℓ , with $\theta \in \{25, 50, 75, 100\}$ and $\ell \in \{1, 2, 3, 4, 5\}$. As shown in Figure 10, ASTRO-MFDF achieves faster convergence on most problems and, in several cases, identifies better solutions, compared to ASTRO-DF and Nelder-Mead, in steady-state simulation optimization.

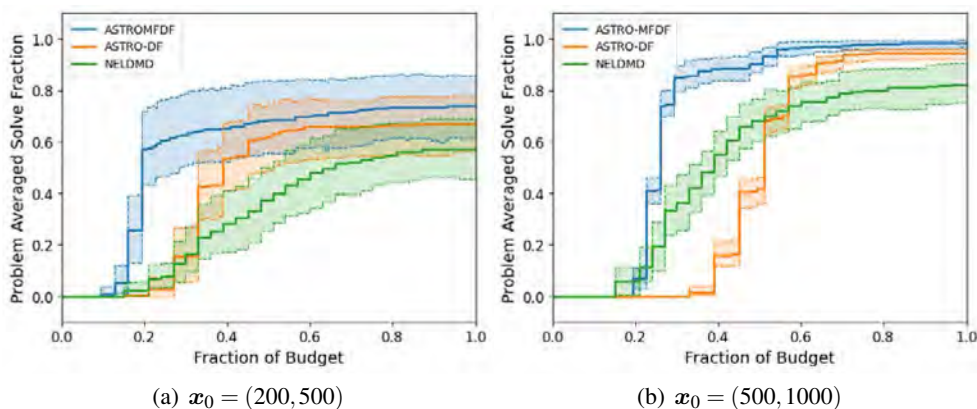


Figure 10: The fraction of problems each method successfully solved under 10% optimality gap threshold.

5 CONCLUSION

In MF methods, information from lower-fidelity models can sometimes hinder, rather than help, the optimization process. Therefore, it is crucial to sequentially determine whether to incorporate lower-fidelity information based on the current state of the optimization. Without this selective approach, using multiple information sources may actually slow convergence. To address this challenge, we introduce a novel stochastic trust-region method, ASTRO-MFDF, designed for MFSO. As a key feature of ASTRO-MFDF, a new adaptive sampling strategy, MFAS, is proposed utilizing MFMC to reduce the variance of function estimates. MFAS dynamically determines the sample sizes for the MF simulations and chooses between MFMC and MC based on sequential estimates of variance and covariance. MFAS helps minimize the waste of computational resources in several ways, such as reducing variance in function estimates and updating correlation vectors when reusing simulation outputs from past runs. Another important feature of ASTRO-MFDF is the incorporation of a correlation vector, which is updated dynamically using information from the optimization history. In particular, if the lower-fidelity models have not contributed to better solutions in past iterations, the method shifts focus to primarily use high-fidelity simulations for the optimization. Through numerical experiments with stochastic Rosenbrock functions and continuous (s, S) inventory problems, we demonstrated that ASTRO-MFDF can achieve faster convergence. This is largely due to its ability to preserve a large trust region for high-fidelity simulations while efficiently leveraging low-fidelity ones, enabling more iterations under a limited budget. For future research, we plan to address high-dimensional and computationally challenging traffic signal control problems in real-world networks by leveraging ASTRO-MFDF with subspace methods and high-performance computing.

ACKNOWLEDGMENTS

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding was provided by the Office of Science, Office of Advanced Scientific Computing

Research, Scientific Discovery through Advanced Computing (SciDAC) program through the FASTMath Institute. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

REFERENCES

- Chang, K.-H., L. J. Hong, and H. Wan. 2013. “Stochastic Trust-Region Response-Surface Method (STRONG)—A New Response-Surface Framework for Simulation Optimization”. *INFORMS Journal on Computing* 25(2):230–243.
- Chen, R., M. Menickelly, and K. Scheinberg. 2018. “Stochastic optimization using a trust-region method and random models”. *Mathematical Programming* 169(2):447–487.
- Chen, X., S. Hemmati, and F. Yang. 2017. “Stochastic co-kriging for steady-state simulation metamodeling”. In *2017 Winter Simulation Conference (WSC)*, edited by W. K. V. Chan, A. D’Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, 1750–1761. IEEE.
- Conn, A. R., K. Scheinberg, and L. N. Vicente. 2009. *Introduction to derivative-free optimization*. 1st ed. Society for Industrial and Applied Mathematics.
- Do, B., and R. Zhang. 2024. “Multi-fidelity Bayesian Optimization: A Review”. *arXiv preprint arXiv:2311.13050*.
- Eckman, D. J., S. G. Henderson, and S. Shashaani. 2023. “SimOpt: A testbed for simulation-optimization experiments”. *INFORMS Journal on Computing* 35(2):495–508.
- D. J. Eckman and S. G. Henderson and S. Shashaani and R. Pasupathy 2023. “SimOpt”. <https://github.com/simopt-admin/simopt>.
- Foumani, Z. Z., A. Yousefpour, M. Shishebor, and R. Bostanabad. 2023. “On the Effects of Heterogeneous Errors on Multi-fidelity Bayesian Optimization”. *arXiv preprint arXiv:2309.02771*.
- Ha, Y., and J. Mueller. 2024. “Adaptive Sampling-Based Bi-Fidelity Stochastic Trust Region Method for Derivative-Free Stochastic Optimization”. *arXiv preprint arXiv:2408.04625*.
- Ha, Y., S. Shashaani, and M. Menickelly. 2025. “Two-stage estimation and variance modeling for latency-constrained variational quantum algorithms”. *INFORMS Journal on Computing* 37(1):125–145.
- Ha, Y., S. Shashaani, and R. Pasupathy. 2024. “Complexity of Zeroth- and First-order Stochastic Trust-Region Algorithms”. *arXiv preprint arXiv:2405.20116*.
- Mainini, L., A. Serani, M. P. Rumpfkeil, E. Minisci, D. Quagliarella, H. Pehlivan, *et al.* 2022. “Analytical benchmark problems for multifidelity optimization methods”. *arXiv preprint arXiv:2204.07867*.
- Nocedal, J., and S. J. Wright. 2006. *Numerical Optimization*. 2nd ed. Springer New York.
- Peherstorfer, B., K. Willcox, and M. Gunzburger. 2016. “Optimal model management for multifidelity Monte Carlo estimation”. *SIAM Journal on Scientific Computing* 38(5):A3163–A3194.
- Sakki, G. K., I. Tsoukalas, P. Kossieris, C. Makropoulos, and A. Efstratiadis. 2022. “Stochastic simulation-optimization framework for the design and assessment of renewable energy systems under uncertainty”. *Renewable and Sustainable Energy Reviews* 168:112886.
- Shashaani, S., F. S. Hashemi, and R. Pasupathy. 2018. “ASTRO-DF: A class of adaptive sampling trust-region algorithms for derivative-free stochastic optimization”. *SIAM Journal on Optimization* 28(4):3145–3176.
- Zhang, Z., Z. Guan, Y. Gong, D. Luo, and L. Yue. 2022. “Improved multi-fidelity simulation-based optimisation: application in a digital twin shop floor”. *International Journal of Production Research* 60(3):1016–1035.

AUTHOR BIOGRAPHIES

YUNSOO HA is a Postdoctoral Researcher at the National Renewable Energy Laboratory within the Artificial Intelligence, Learning, and Intelligent Systems (ALIS) group. Prior to this, he earned Ph.D. from the Department of Industrial and Systems Engineering at North Carolina State University. His research revolves around machine learning, quantum computing, and stochastic optimization. His email address is yunsoo.ha@nrel.gov.

JULIANE MUELLER is the manager of the Artificial Intelligence, Learning, and Intelligent Systems (ALIS) group within the Computational Science Center at NREL. Her background is in the development of numerical optimization algorithms for black-box and compute-intensive problems where analytic descriptions of objective and constraint functions are not available. Her algorithm developments include surrogate modeling, machine and active learning with application to various DOE problems. Her email address is juliane.mueller@nrel.gov.