IISE Transactions

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/uiie21

Iteration complexity and finite-time efficiency of adaptive sampling trust-region methods for stochastic derivative-free optimization

Yunsoo Ha & Sara Shashaani

To cite this article: Yunsoo Ha & Sara Shashaani (04 Jun 2024): Iteration complexity and finite-time efficiency of adaptive sampling trust-region methods for stochastic derivative-free optimization, IISE Transactions, DOI: 10.1080/24725854.2024.2335513

To link to this article: https://doi.org/10.1080/24725854.2024.2335513

IISE TRANSACTIONS
https://doi.org/10.1080/24725854.2024.2335513

Iteration complexity and finite-time efficiency of adaptive sampling trust-region methods for stochastic derivative-free optimization

Yunsoo Ha and Sara Shashaani

Edward P. Fitts Department of Industrial and System Engineering, North Carolina State University, Raleigh, NC, USA

ABSTRACT

ASTRO-DF is a prominent trust-region method using *adaptive sampling* for stochastic derivative-free optimization of nonconvex problems. Its salient feature is an easy-to-understand-and-implement concept of maintaining "just enough" replications when evaluating points throughout the search to guarantee almost-sure convergence to a first-order critical point. To reduce the dependence of ASTRO-DF on the problem dimension and boost its performance in finite time, we present two key refinements, namely: (i) local models with diagonal Hessians constructed on interpolation points based on a coordinate basis; and (ii) direct search using the interpolation points whenever possible. We demonstrate that the refinements in (i) and (ii) retain the convergence guarantees while matching existing results on iteration complexity. Uniquely, our iteration complexity results match the canonical rates without placing assumptions on iterative models' quality and their independence from function estimates. Numerical experimentation on a testbed of problems and comparison against existing popular algorithms reveals the computational advantage of ASTRO-DF due to the proposed refinements.

ARTICLE HISTORY
Received 12 June 2023
Accepted 19 March 2024

KEYWORDS
Simulation optimization; zeroth-order oracle; finite-time performance

1. Introduction

We consider unconstrained Simulation Optimization (SO) problems of the form

$$\min_{\boldsymbol{x}\in\mathbb{R}^d} f(\boldsymbol{x}) := \mathbb{E}[F(\boldsymbol{x},\xi)] = \int_{\Xi} F(\boldsymbol{x},\xi)dP(\xi), \tag{1}$$

where $f:\mathbb{R}^d\to\mathbb{R}$ is smooth (but nonconvex) and bounded below, and $F:\mathbb{R}^d\times\Xi\to\mathbb{R}$ is the stochastic function value defined on the probability space $(\Xi,\mathcal{F},P)$. In particular, consider $f(\boldsymbol{x})$ that is only observable with noise by a Monte Carlo simulation, which generates the random variable $F(\boldsymbol{x},\xi)$, that can be decomposed into the true function value and the stochastic error $E(\boldsymbol{x},\xi)$, i.e., $F(\boldsymbol{x},\xi)=f(x)+E(\boldsymbol{x},\xi)$. The estimator of $f(\boldsymbol{x})$ via sample average approximation and the estimated variance of $F(\boldsymbol{x},\xi)$ with n runs of the simulation can be obtained by

$$\bar{F}(\boldsymbol{x},n) = n^{-1}\sum_{i=1}^{n} F(\boldsymbol{x},\xi_i), \text{ and}$$

$$\hat{\sigma}_F^2(\boldsymbol{x},n) = (n-1)^{-1}\sum_{j=1}^{n}\left(F(\boldsymbol{x},\xi_i)-\bar{F}(\boldsymbol{x},n)\right)^2.$$

If derivative information is not directly available from the Monte Carlo Simulation, then Problem (1) becomes a Stochastic Derivative-Free Optimization (SDFO) problem. A SDFO algorithm produces, in one run, $\{\boldsymbol{X}_k\}$—a sequence of stochastic incumbents at each iteration $k\in\mathbb{N}$. This sequence can be viewed as a stochastic process defined on a filtered probability space $(\Omega,\mathcal{F}_k,\mathbb{P})$, where $\mathcal{F}_k$ denotes the σ-algebra increasing in k. The goal is designing an efficient SDFO algorithm to reach an ϵ-stationary point, that is a point in the set $\{\boldsymbol{x}:\|\nabla f(\boldsymbol{x})\|\le\epsilon\}$, with probability 1.

1.1. Adaptive sampling and trust regions

The number of function evaluations n at an incumbent $\boldsymbol{X}_k$ generated by the SDFO algorithm must guarantee estimation accuracy of

$$\mathbb{P}\{|\bar{E}(\boldsymbol{X}_k,n)|>\epsilon_k\}\le\alpha_k, \tag{2}$$

with the required accuracy threshold $\epsilon_k>0$ and exceedance probability $\alpha_k\in(0,1)$, where $\bar{E}(\boldsymbol{X}_k,n):=\bar{F}(\boldsymbol{X}_k,n)-f(\boldsymbol{X}_k)$ denotes the estimation error. An insufficiently large n during the optimization will threaten the convergence of the SDFO algorithm whereas an unnecessarily large n at every design point leads to exhausting the simulation budget before reaching a good solution. Therefore, choosing sufficiently large n for each visited design point, that is, viewing the sample size as a function of $\boldsymbol{X}_k$ denoted by $n(\boldsymbol{X}_k)$ not only secures strong consistency, but also reinforces efficiency in finite-time and asymptotically.

There is emerging evidence (Gratton *et al.*, 2018; Maggiar *et al.*, 2018) that stochastic trust-region methods are effective at solving nonconvex SDFO, mainly due to their natural ability to self-tune step sizes and facility for curvature estimation. ASTRO-DF—Adaptive Sampling Trust-region

CONTACT Sara Shashaani sshasha2@ncsu.edu
Supplemental data for this article is available online at https://doi.org/10.1080/24725854.2024.2335513.

Optimization Derivative-free—(Shashaani *et al.*, 2016, 2018; Ha *et al.*, 2021) is a class of trust-region optimization algorithms suited for SDFO that deals with the challenge of choosing the right sample size using an adaptive sampling strategy. The sample size is decided on the fly as a stopping time random variable. For a visited incumbent $\boldsymbol{X}_k$, using a random sample size $N(\boldsymbol{X}_k)$ that adapts to ϵ_k and α_k at iteration k in (2) based on the inferred estimation error from the collected observations can lead to convergence to a first-order stationary point almost surely. Given its versatile mechanics, guarantees in theory with few requirements, and promising results from straightforward implementation, capitalizing on adding new features to ASTRO-DF to make it more flexible for high-dimensional SDFO problems is favorable. To this end, the current article proposes refinements to the algorithm and establishes theoretical and finite-time efficiency properties under those refinements.

The progress in ASTRO-DF relies on a local model that, due to the derivative-free setting, is constructed on function value estimates of points near the incumbent within a trust region of size Δ_k that implicitly controls the step size. The mechanism of derivative-free trust-region optimization dictates that Δ_k eventually drops to zero to drive the algorithm's convergence. Approximately minimizing the local model within the trust region (a.k.a., the subproblem) provides a candidate for the next incumbent. The decreasing accuracy threshold for each point in iteration k, ϵ_k, is designed to decrease linearly with Δ_k^2 and the exceedance probability α_k drops to zero in k due to the assurance that $N(\boldsymbol{X}_k)$ diverges almost surely. In fact, $N(\boldsymbol{X}_k)$ increases linearly with Δ_k^{-4} as Δ_k decays to zero.

1.2. Summary of insights and contributions

Besides the rapid increase in sample size, which is common in SDFO, ASTRO-DF is vulnerable to some practical disadvantages:

1. $\mathcal{O}(d^2)$ number of points needed to build a quadratic model to capture the local curvature information is too costly, especially considering that an increasing number of oracle runs is required in each point with k. Furthermore, random choice of the design points causes practical inefficiency: to obtain a good model quality, the design points need to be "well-poised" (or well spread) in the trust region; this entails linear algebra cost of the order d^6 for a quadratic model; see Chapter 6 by Conn *et al.* (2009).
2. Even though ASTRO-DF converges to a stationary point almost surely, the probability of having a successful iteration, which is the probability of finding a new incumbent at each iteration, can start off very small, tending to one as $k \to \infty$. As a practical matter, frequently rejecting the candidate incumbent suggested by the model reduces the trust-region radius Δ_k too quickly, which in turn explodes the sample size leading to exhausted simulation budget before finding a good solution. In other words, full reliance on the local model can slow the algorithm's progress.

To address these challenges, we propose two refinements. The first refinement is to use a fixed geometry for the design set based on the coordinate basis leading to a model with a diagonal Hessian (Coope and Tappenden, 2021). Such a fixed geometry was recently shown to be optimal among all designs of $2d+1$ many points (Ragonneau and Zhang, 2023). Selecting the points with this design, in addition to eliminating a source of randomness inside the algorithm, also eliminates the linear algebra cost of certifying that the design is well-poised and the row operations needed to solve the system of equations in the interpolation to construct the model. More importantly, the order of approximated gradient's deterministic error (bias) is reduced to $\mathcal{O}(\Delta_k^2)$, that is a central-difference like error, from $\mathcal{O}(\Delta_k)$, that is a forward-difference like error, in the original algorithm (see Theorem 4.4).

The second refinement is to leverage the design points in a direct search manner. This refinement is motivated by observing that in each iteration of ASTRO-DF, the estimated function value at the candidate incumbent recommended by the local model can be worse than that of the points used to construct the model. Despite such a possibility, the original setup deems the iteration unsuccessful if the model candidate yields insufficient reduction, ignoring the readily identified points that obtain improvement, and starts over to build a model in a smaller neighborhood. In addition to smaller steps, this rejection will increase oracle calls, incurring a considerable waste of simulation budget. A direct search strategy considers selecting design points whose estimated function values provide sufficient reduction in the function value estimates as the next incumbent, enabling progress without extra cost and maintaining stability in step size and sample size in the early iterations.

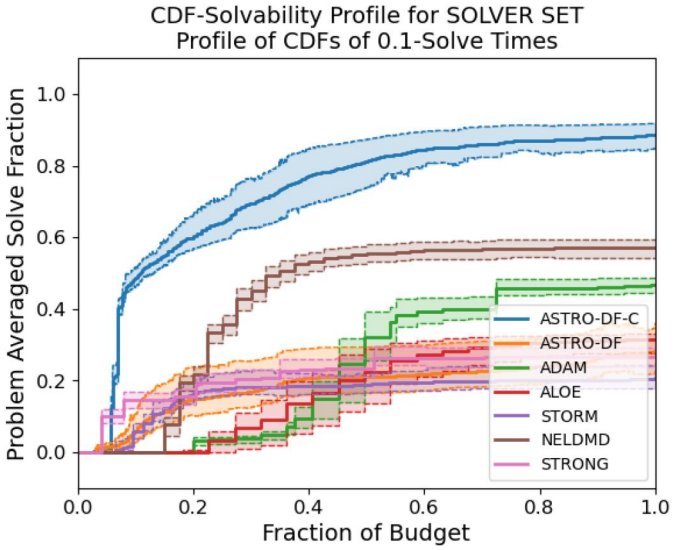

Figure 1. Fraction of 60 problems from SimOpt library (Eckman *et al.*, 2021; Eckman *et al.*, 2023) solved to 0.1-optimality with 95% confidence intervals from 20 runs of each algorithm shows a clear advantage in finite-time performance of ASTRO-DF-C.

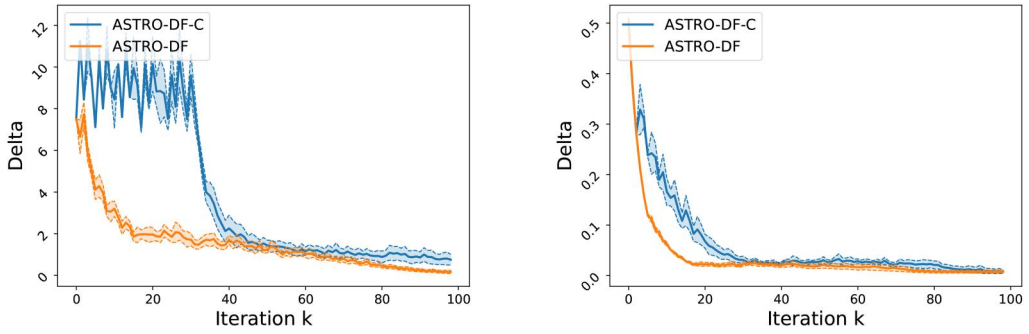

(a) 20-D Noisy Rosenbrock (b) 13-D Stochastic Activity Network

Figure 2. Slower rate of trust-region radius decay in the early iterations of the search as a result of using direct search allows for larger step size, and hence, better progress. The mean and 95% confidence interval of trajectory of Δ_k exhibits this behavior for two problems.

We name ASTRO-DF with coordinate direct search refinements as ASTRO-DF-C. These refinements lead to significantly better performance (see Figure 1) compared with the state-of-the-art SDFO algorithms. As exhibited in Figure 2, the boost in finite-time performance of ASTRO-DF-C is intuitive: direct search increases the probability of finding a new incumbent and prevents the trust-region size from becoming too small (see Section 5 for improvements in terms of objective function values). As $k \to \infty$, the success probability will match that in ASTRO-DF, since the models become progressively better approximations of the function and are more likely to suggest better incumbents than the points identified with the coordinate basis.

1.3. Organization of the article

After a review of the existing body of work for SDFO in Section 2, we detail ASTRO-DF-C in Section 3. Importantly, and as we show in Section 4, gains in finite-time performance are realized without compromising guarantees on convergence. We augment ASTRO-DF-C's convergence results with proof of ϵ^{-2} order of the number of iterations to reach ϵ-optimality in expectation. This result is obtained despite the relaxed requirements on quality and independence of local models from function estimates that was a necessity to prove iteration complexity in earlier work (Blanchet *et al.*, 2019). Section 5 explores extensive numerical experiments on a convex and a nonconvex problem, both considered expensive to solve in the derivative-free setting. We will summarize our findings and future needs in Section 6.

2. Preliminaries

In this section, we introduce notations used in the article, review the existing literature for SDFO, and provide essential definitions.

2.1. Notation and terminology

We use bold font for vectors; $\boldsymbol{x} = (x_1, x_2, ..., x_d) \in \mathbb{R}^d$ denotes a d-dimensional vector of real numbers. Let $\boldsymbol{e}^i \in \mathbb{R}^d$ for $i = 1, ..., d$ denote the standard unit basis vector in $\mathbb{R}^d$. We use calligraphic fonts for sets and sans serif fonts for matrices. $\mathcal{A}^{\complement}$ denotes the complement of the set $\mathcal{A}$. $\mathcal{B}(\boldsymbol{x}; \Delta) = \{\boldsymbol{x} + \boldsymbol{s} \in \mathbb{R}^d : \|\boldsymbol{s}\|_2 \leq \Delta\}$ denotes the closed ball of radius $\Delta > 0$ with center $\boldsymbol{x}$. For a sequence of sets $\{\mathcal{A}_n\}$, the set $\{\mathcal{A}_n \text{ i.o.}\}$ denotes $\text{limsup}_{n\to\infty} \mathcal{A}_n$. For two functions $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$, we say $f(\boldsymbol{x}) = \mathcal{O}(g(\boldsymbol{x}))$ if there exist positive numbers ϵ and M such that $|f(\boldsymbol{x})| \leq Mg(\boldsymbol{x})$ for all $\boldsymbol{x}$ with $|\boldsymbol{x}| < \epsilon$. We use $a \wedge b := \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. We use capital letters for random scalars and vectors. For a sequence of random vectors $\{\boldsymbol{X}_k\}, k \in \mathbb{N}$, $\boldsymbol{X}_k \overset{w.p.1}{\to} \boldsymbol{X}$ denotes convergence almost surely as $k \to \infty$. We say $X_n = \mathcal{O}_p(1)$ when the random sequence $\{X_n\}$ is stochastically bounded, i.e., given $\epsilon > 0$ there exist $n_\epsilon \in \mathbb{N}$ and $c_\epsilon > 0$ such that $\mathbb{P}\{|X_n| > c_\epsilon\} < \epsilon$ for all $n \geq n_\epsilon$. A random function $F : \mathbb{R}^d \times \Xi \to \mathbb{R}$ defined on the probability space $(\Xi, \mathcal{F}, \mathbb{P})$ is *strongly* uniformly bounded if a constant $\kappa_{Eb} \in \mathbb{R}$ bounds $F(\boldsymbol{x}, \xi)$ with probability 1 for all $\boldsymbol{x} \in \mathbb{R}^d$. The term "iid", abbreviates independent and identically distributed.

2.2. Literature review

In the artificial intelligence era, derivative-free optimization has received much attention for allowing the user to specify the objective function involved in non-explicit forms. As a result, derivative-free optimization has a wide range of applications such as hyper-parameter tuning (Ghanbari and Scheinberg, 2017; Ruan *et al.*, 2020; Cakmak *et al.*, 2021), reinforcement learning (Flaxman *et al.*, 2004; Salimans *et al.*, 2017; Choromanski *et al.*, 2018; Fazel *et al.*, 2018), simulation-based optimization (Chang *et al.*, 2013; Xu and Nelson, 2013), game theory (Zhang *et al.*, 2023), and quantum computing (Menickelly *et al.*, 2023). An essential characteristic of SDFO is that the function evaluations are only accessible via a black-box simulation with stochastic error. Running this noisy simulation can be expensive. Hence, one of the persistent aims in this field is to improve the efficiency of the algorithms (Shashaani *et al.*, 2018; Paquette and Scheinberg, 2020; Jin *et al.*, 2021).

Efficiency guarantees of an algorithm for SDFO are secondary to the convergence guarantees that imply the iterates produced by the algorithm stabilize towards a stationary point in some probabilistic sense. An adaptive sampling

strategy, which adapts the number of required function evaluations at each design point to the inferred closeness to optimality, leads to almost sure convergence to a stationary point. The intuition behind adaptive sampling is that more accurate estimates are needed in near-optimal regions to guarantee we can identify better incumbents.

In the existing literature, the efficiency of SDFO algorithms is typically evaluated through a complexity analysis that assesses the costs required to achieve a near-optimal solution. The costs can take different forms, such as the number of iterations known as "iteration complexity" or number of function evaluations or simulation runs known as "work complexity". Work complexity is a more appropriate metric for gauging the computational load of SDFO algorithms, as the number of function evaluations could vary per iteration for the almost sure convergence. The existing body of work that analyzes SDFO algorithms has little known work complexity guarantees but has focused on attaining the canonical rate of expected iteration complexity $\mathbb{E}[T_\epsilon] = \mathcal{O}(\epsilon^{-2})$, where

$$T_\epsilon = \{k \in \mathbb{N} : \|\nabla f(\boldsymbol{X}_k)\| \leq \epsilon\}$$

is the first iteration reaching ϵ-optimality. The complexity of SDFO problems proliferates with the problem dimension, due to the necessity of the derivative estimator using only expensive, noisy simulations. The method for the derivative approximation and the sample size of each design point are the key factors in the complexity and finite-time performance. The two suggested refinements for ASTRO-DF here aim for solving higher dimensional derivative-free problems more efficiently.

Recall, an algorithm for SDFO needs function estimates, and the accuracy of the estimates is important for convergence. Hence, we review the existing literature based on assumptions on the stochastic error at $\boldsymbol{x} \in \mathbb{R}^d$, namely $E(\boldsymbol{x}, \xi)$, whose expectation may or may not be zero. Importantly, in our literature review we primarily focus on algorithms that are devised for a nonconvex function f of the form of an expected value, as specified in (1). Although the body of literature for convex settings is vast, we find two recent studies (Hu *et al.*, 2023; Xu and Zheng, 2023) that address convex settings for simulation optimization with derivative-free algorithms worthy of note. We also exclude literature on optimization over discrete-space variables. The remainder of this section returns to non-convex stochastic optimization.

2.2.1. Strong uniform boundedness of $|E(x, \xi)|$

Several methods for approximating gradients, such as a finite difference and Gaussian smoothing, are analyzed with bounds for the number of design points to obtain sufficiently accurate estimates (Berahas Cao, Choromanski and Scheinberg, 2021). The expected iteration complexity of $\mathcal{O}(\epsilon^{-2})$ is obtained in the generic line search that converges to an L_2-ball centered around a stationary point with a radius depending on the uniform upper bound on $|E(\boldsymbol{x}, \xi)|$ (Berahas Cao and Scheinberg, 2021). This *optimal neighborhood* can be large in more noisy settings impacting the ability to attain a satisfactory solution.

2.2.2. One-sided sub-exponential $E(x, \xi)$

The Adaptive Line search with Oracle Estimations algorithm (ALOE) (Jin *et al.*, 2021) in this setting converges to an optimal neighborhood with high probability tail bound for the iteration complexity $\mathbb{P}\{T_\epsilon \leq t\} \geq 1 - \mathcal{O}(e^{-t})$. The size of the optimal neighborhood depends on the upper bound of $\mathbb{E}[E(\boldsymbol{x}, \xi)]$. The trust-region method with this assumption on the stochastic errors has recently been developed (Cao *et al.*, 2022) exhibiting the same iteration complexity. One-sided sub-exponentiality is a weaker assumption than sub-exponentiality of $E(\boldsymbol{x}, \xi)$ as it only requires for the stochastic error not to be too large in the sense that $\mathbb{E}$ i.e., $\mathbb{E}\left[e^{-\lambda(E(\boldsymbol{x},\xi)-\mathbb{E}[E(\boldsymbol{x},\xi)])}\right] \leq e^{\lambda\sigma/2}$ for some $\sigma > 0$ and some small enough $\lambda > 0$.

2.2.3. Zero-mean finite-variance $E(x, \xi)$

The Trust Reign Optimization with Random Models (STORM) algorithm (Chen *et al.*, 2018) using this assumption converges to a stationary point almost surely with the expected iteration complexity $\mathcal{O}(\epsilon^{-2})$ (Blanchet *et al.*, 2019). The convergence and complexity analysis for STORM relies on sufficiently replicating the function value at any given point $\boldsymbol{x}$ such that given $\epsilon > 0$, $\mathbb{P}\{|\bar{E}(\boldsymbol{X}_k, N(\boldsymbol{X}_k))| > \epsilon\} \leq \alpha$ for small enough α. Moreover, in order to meet the condition for any $k \in \mathbb{N}$, STORM necessitates an extra assumption, namely that the function estimates and local models are independent. This implies that the information obtained in previous iterations cannot be reused in the current iteration, resulting in inefficiencies. ASTRO-DF assumes zero-mean stochastic errors with bounded νth moment for $2 \leq \nu \leq 8$; hence, ensuring finite variance. Under these assumptions, ASTRO-DF converges to a stationary point almost surely following $\mathbb{P}\{|\bar{E}(\boldsymbol{X}_k, N(\boldsymbol{X}_k))| > \epsilon\} \leq \alpha_k$, where $\alpha_k \to 0$ as $k \to \infty$.

2.2.4. Decaying $\mathbb{E}[|E(x, \xi)|]$ with distance to stationarity and bounded variance

A backtracking Armijo line search method (Paquette and Scheinberg, 2020) achieves $\mathcal{O}(\epsilon^{-2})$ as the expected iteration complexity if the bias converges to zero as $(\Delta_k^4 \vee \|\nabla f(\boldsymbol{X}_k)\|^4) \to 0$ while $\mathrm{Var}(E(\boldsymbol{x}, \xi))$ (in addition to the variance of the gradient error if its observations are directly available) is finitely bounded for all $\boldsymbol{x}$. The function and gradient estimates are also assumed to be accurate with a high probability as in STORM.

In contrast with the related work in the literature, our assumptions in this article for the nature of stochastic error are that it has zero mean and exhibits a sub-exponential tail for all $\boldsymbol{x} \in \mathrm{IR}^d$. We formalize this in Assumption 2. The sub-exponential family includes random variables with heavy tails, such as those with lognormal distribution, that do not necessarily have finite moment-generating functions. In addition, no restriction on the expected absolute value (absolute integrability) is required. This implies that our framework and results are more general and pertain to a larger range of stochastic systems. Although we assume zero-mean stochastic error, the roadmap to an extension of

our work with biased random observations that converges to a neighborhood of stationarity seems evident and left for future research.

2.3. Definitions

We now introduce several definitions in all of which, we start with a center point $\boldsymbol{x}$ (later, the incumbent $\boldsymbol{X}_k$) and the trust-region radius Δ (later, Δ_k that changes with k).

Definition 2.1. *(stochastic polynomial interpolation models) Given $\boldsymbol{x} \in \mathbb{R}^d$ and $\Delta > 0$, let $\Phi(\boldsymbol{x}) = (\phi_0(\boldsymbol{x}), \phi_1(\boldsymbol{x}), ..., \phi_q(\boldsymbol{x}))$ be a polynomial basis on $\mathbb{R}^d$. With $p = q$ and the design set $\mathcal{X} := \{\boldsymbol{x}^0, \boldsymbol{x}^1, ..., \boldsymbol{x}^p\} \subset \mathcal{B}(\boldsymbol{x}; \Delta)$, we can find $\boldsymbol{\beta} = (\beta_0, \beta_1, ..., \beta_p)$ such that $\mathsf{M}(\Phi, \mathcal{X})\boldsymbol{\beta} = \bar{F}(\mathcal{X}, N(\mathcal{X}))$, where*

$$\mathsf{M}(\Phi, \mathcal{X}) = \begin{bmatrix} \phi_1(\boldsymbol{x}^0) & \phi_2(\boldsymbol{x}^0) & \cdots & \phi_q(\boldsymbol{x}^0) \\ \phi_1(\boldsymbol{x}^1) & \phi_2(\boldsymbol{x}^1) & \cdots & \phi_q(\boldsymbol{x}^1) \\ \vdots & \vdots & \vdots & \vdots \\ \phi_1(\boldsymbol{x}^p) & \phi_2(\boldsymbol{x}^p) & \cdots & \phi_q(\boldsymbol{x}^p) \end{bmatrix},$$

$$\bar{F}(\mathcal{X}, N(\mathcal{X})) = \begin{bmatrix} \bar{F}(\boldsymbol{x}^0, N(\boldsymbol{x}^0)) \\ \bar{F}(\boldsymbol{x}^1, N(\boldsymbol{x}^1)) \\ \vdots \\ \bar{F}(\boldsymbol{x}^p, N(\boldsymbol{x}^p)) \end{bmatrix},$$

assuming that the so-called Vandermonde matrix $\mathsf{M}(\Phi, \mathcal{X})$ is invertible. We note $\boldsymbol{x}^0 := \boldsymbol{x}$. For the matrix $\mathsf{M}(\Phi, \mathcal{X})$ to be nonsingular, the set $\mathcal{X}$ should be poised in $\mathcal{B}(\boldsymbol{x}; \Delta)$. Then, the function $M : \mathcal{B}(\boldsymbol{x}; \Delta) \to \mathbb{R}$, defined as $M(\boldsymbol{x}) = \sum_{j=0}^{p} \beta_j \phi_j(\boldsymbol{x})$ is a stochastic polynomial interpolation model of f on $\mathcal{B}(\boldsymbol{x}; \Delta)$. Let $\boldsymbol{G} = [\beta_1 \beta_2 \cdots \beta_d]^\top$ be the subvector of $\boldsymbol{\beta}$ and H be a symmetric matrix of size $d \times d$ with elements uniquely defined by $\begin{bmatrix} \beta_{d+1} & \beta_{d+2} & \cdots & \beta_p \end{bmatrix}^\top$. Then, the stochastic quadratic model $M : \mathcal{B}(\boldsymbol{x}; \Delta) \to \mathbb{R}$, is

$$M(\boldsymbol{x} + \boldsymbol{s}) = \beta_0 + \boldsymbol{s}^\top \boldsymbol{G} + \frac{1}{2}\boldsymbol{s}^\top \mathsf{H}\boldsymbol{s}. \tag{3}$$

Definition 2.2. *(stochastic quadratic models with diagonal Hessians) A special case of (3) is when the Hessian has only diagonal values, i.e.,*

$$\mathsf{H} = \begin{bmatrix} H_1 & & 0 \\ & \ddots & \\ 0 & & H_d \end{bmatrix} \in \mathbb{R}^{d \times d}. \tag{4}$$

In the stochastic quadratic interpolation model with diagonal Hessian, $p = 2d$, the model (3) contains $2d + 1$ unknowns, and $2d + 1$ function value estimations are needed to uniquely determine the $\boldsymbol{G}$ and H, letting the interpolation set be

$$\mathcal{X}_{cb} = \{\boldsymbol{x}^0, \boldsymbol{x}^0 + \boldsymbol{e}^1\Delta, ..., \boldsymbol{x}^0 + \boldsymbol{e}^d\Delta, \boldsymbol{x}^0 - \boldsymbol{e}^1\Delta, ..., \boldsymbol{x}^0 - \boldsymbol{e}^d\Delta\}$$

contained in $\mathcal{B}(\boldsymbol{x}^0; \Delta)$. Since the coordinate basis is used to generate the interpolation set, $\boldsymbol{\beta}$ is guaranteed to exist. Hence, $H_i = \beta_{d+i} < \infty$ for all $i = 1, 2, ..., d$. In this case, $\Phi(\boldsymbol{x}) := (1, x_1, x_2, ..., x_d, x_1^2, x_2^2, ..., x_d^2)$, and $M(\cdot)$ is said to be a stochastic quadratic model with diagonal Hessian.

Definition 2.3. *(stochastic fully linear models) Given $\boldsymbol{x} \in \mathbb{R}^d$ and $\Delta > 0$, a function $M : \mathcal{B}(\boldsymbol{x}; \Delta) \to \mathbb{R}$ obtained following Definition 2.2 is a stochastic fully linear model of f on $\mathcal{B}(\boldsymbol{X}; \Delta)$ if ∇f is Lipschitz continuous with constant κ_L, and there exist positive constants κ_{eg} and κ_{ef} dependent on κ_L but independent of $\boldsymbol{x}$ and Δ such that almost surely*

$$\|\nabla f(\boldsymbol{x}) - \nabla M(\boldsymbol{x})\| \le \kappa_{eg}\Delta \text{ and}$$
$$|f(\boldsymbol{x}) - M(\boldsymbol{x})| \le \kappa_{ef}\Delta^2 \forall \boldsymbol{x} \in \mathcal{B}(\boldsymbol{x}; \Delta).$$

Definition 2.4. *(Cauchy reduction) Given $\boldsymbol{x} \in \mathbb{R}^d$, $\Delta > 0$, and a function $M : \mathcal{B}(\boldsymbol{x}; \Delta) \to \mathbb{R}$ obtained following Definition 2.2, $\boldsymbol{s}^c$ is called the Cauchy step if*

$$M(\boldsymbol{x}) - M(\boldsymbol{x} + \boldsymbol{s}^c) \ge \frac{1}{2}\|\nabla M(\boldsymbol{x})\| \left(\frac{\|\nabla M(\boldsymbol{x})\|}{\|\nabla^2 M(\boldsymbol{x})\|} \wedge \Delta \right). \tag{5}$$

We assume that $\|\nabla M(\boldsymbol{x})\|/\|\nabla^2 M(\boldsymbol{x})\| = +\infty$ when $\|\nabla^2 M(\boldsymbol{x})\| = 0$. The Cauchy step is obtained by minimizing the model $M(\cdot)$ along the steepest descent direction within $\mathcal{B}(\boldsymbol{x}; \Delta)$, and hence, easy and quick to obtain.

Definition 2.5. *(filtration and stopping time) A filtration $\{\mathcal{F}_k\}_{k \ge 1}$ over a probability space $(\Omega, \mathbb{P}, \mathcal{F})$ is defined as an increasing family of σ-algebras of $\mathcal{F}$, i.e., $\mathcal{F}_k \subset \mathcal{F}_{k+1} \subset \mathcal{F}$ for all k. We interpret $\mathcal{F}_k$ as "all the information available at time k." A filtered space $(\Omega, \mathbb{P}, \{\mathcal{F}_k\}_{k \ge 1}, \mathcal{F})$ is a probability space equipped with a filtration. A map $N : \Omega \to \{0, 1, 2, ..., \infty\}$ is called a stopping time with respect to filtration $\mathcal{F}$ if the event $\{N = n\} := \{\omega : N(\omega) = n\} \in \mathcal{F}$, i.e., it is $\mathcal{F}$-measurable for all $n < \infty$.*

3. ASTRO-DF with coordinate direct search

Let us first briefly review ASTRO-DF. Each iteration of ASTRO-DF consists of four steps:

1. Constructing a local model by interpolating function value estimates (with adaptive sample sizes) on a randomly selected *poised* (or well spread) design set and repeatedly shrinking the trust region and updating the model if the model gradient is too small relative to the trust-region size.
2. Recommending a candidate incumbent that approximately minimizes the model.
3. Evaluating the reduction in the function value after estimating it at the candidate incumbent.
4. Updating the trust-region radius and incumbent, i.e., accepting the candidate and expanding the trust-region radius if reduction in the function value is sufficient (relative to the model reduction) or rejecting the candidate and shrinking the trust-region radius otherwise. See (Shashaani *et al.*, 2018) for more extensive details.

The refinements in ASTRO-DF-C are in the first and last step as well as in the sample size lower bound for the simulation budget. In step 1, instead of random selection of the design set that would also require ensuring its well-poised geometry in the trust region, the algorithm selects interpolation points along the coordinate basis. The repeated shrinkage and updating, which is known as the *criticality step* in the DFO literature (Conn *et al.*, 2009), is removed as an inner loop and transferred to step 4 and acceptance criteria. In step 4, accepting a new incumbent is augmented using a direct search strategy. The details of these refinements and their consequences are described below and the complete steps of ASTRO-DF-C are listed in Algorithm 1.

Algorithm 1 ASTRO-DF-C: ASTRO-DF with Coordinate Direct Search

Required: Initial guess $\boldsymbol{x}_0 \in \mathbb{R}^d$, initial and maximum trust-region radius $\Delta_0, \Delta_{\max} > 0$, model "fitness" and certification thresholds $\eta \in (0,1)$ and $\mu > 0$, sufficient reduction constant $\theta > 0$, expansion and shrinkage constants $\gamma_1 > 1$ and $\gamma_2 \in (0,1)$, sample size lower bound λ_k, and adaptive sampling constant $\kappa > 0$.

1: **for** $k = 0, 1, 2, \ldots$ **do**

2: *Model Construction:* Select the design set $\mathcal{X}_k = \{\boldsymbol{X}_k^i\}_{i=0}^{2d} \subset \mathcal{B}(\boldsymbol{X}_k; \Delta_k)$ following Equation (3) and estimate $\bar{F}(\boldsymbol{X}_k^i, N(\boldsymbol{X}_k^i))$, where

$$N(\boldsymbol{X}_k^i) = \min\left\{ n \geq \lambda_k : \frac{\hat{\sigma}(\boldsymbol{X}_k^i, n)}{\sqrt{n}} \leq \frac{\kappa \Delta_k^2}{\sqrt{\lambda_k}} \right\}, \tag{6}$$

for $i = 0, 1, \ldots, 2d$ ($\boldsymbol{X}_k^0 = \boldsymbol{X}_k$) and construct the model $M_k(\boldsymbol{X}_k + \boldsymbol{s})$ via interpolation.

3: *Subproblem:* Approximate the kth step by minimizing the model in the trust region, $\boldsymbol{S}_k = \operatorname{argmin}_{\|\boldsymbol{s}\| \leq \Delta_k} M_k(\boldsymbol{X}_k + \boldsymbol{s})$, and set $\tilde{\boldsymbol{X}}_{k+1} = \boldsymbol{X}_k + \boldsymbol{S}_k$.

4: *Candidate Evaluation:* Use adaptive sampling rule (6) to estimate $\bar{F}(\tilde{\boldsymbol{X}}_{k+1}, \tilde{N}_{k+1})$. Define the best design point $\hat{\boldsymbol{X}}_{k+1} = \operatorname{argmin}_{\boldsymbol{x} \in \mathcal{X}_k \setminus \boldsymbol{X}_k} \bar{F}(\boldsymbol{x}, N(\boldsymbol{x}))$, its sample size $\hat{N}_{k+1} = N(\hat{\boldsymbol{X}}_{k+1})$, incumbent's sample size $\hat{N}_k = N(\boldsymbol{X}_k)$, direct search reduction $\hat{R}_k = \bar{F}(\boldsymbol{X}_k, \hat{N}_k) - \bar{F}(\hat{\boldsymbol{X}}_{k+1}, \hat{N}_{k+1})$, subproblem reduction $\tilde{R}_k = \bar{F}(\boldsymbol{X}_k, \hat{N}_k) - \bar{F}(\tilde{\boldsymbol{X}}_{k+1}, \tilde{N}_{k+1})$, and model reduction $R_k = M_k(\boldsymbol{X}_k) - M_k(\tilde{\boldsymbol{X}}_{k+1})$.

5: *Update:* Set

$$(\boldsymbol{X}_{k+1}, N_{k+1}, \Delta_{k+1}) = \begin{cases} (\hat{\boldsymbol{X}}_{k+1}, \hat{N}_{k+1}, \gamma_1 \Delta_k \wedge \Delta_{\max}) & \text{if } \hat{R}_k > (\tilde{R}_k \vee \theta \Delta_k^2), \\ (\tilde{\boldsymbol{X}}_{k+1}, \tilde{N}_{k+1}, \gamma_1 \Delta_k \wedge \Delta_{\max}) & \text{else if } \tilde{R}_k \geq \eta R_k \text{ and } \mu \|\nabla M_k(\boldsymbol{X}_k)\| \geq \Delta_k, \\ (\boldsymbol{X}_k, \hat{N}_k, \Delta_k \gamma_2) & \text{otherwise,} \end{cases}$$

and $k = k + 1$.

6: **end for**

3.1. Model construction

In ASTRO-DF-C, the point selection is deterministic and the local model $M_k(\cdot)$ is constructed with $2d$ new points along the coordinate basis at the trust-region boundary in each iteration to obtain a diagonal Hessian model. As mentioned in Section 1, the suggested design set has several advantages. First, the model gradient error for all points within a trust region has a similar upper bound, but with fewer design points. We will also show in Theorem 4.4 that the deterministic part of the gradient error (bias) at the center point improves to $\mathcal{O}(\Delta_k^2)$ from $\mathcal{O}(\Delta_k)$. Second, the design set via coordinate basis is optimal in the sense that it minimizes the constant of well-poisedness in a ball, saving the linear algebra cost to make a well-poised set (Ragonneau and Zhang, 2023). Since the coordinate basis automatically provides a well-poised geometry, the computational cost of ensuring well-poisedness, which also depends on the problem dimension, e.g., $\mathcal{O}(d^2)$ operations for a linear model and $\mathcal{O}(d^6)$ for a quadratic model (Conn *et al.*, 2009) [Algorithms 6.1-6.3], is completely eliminated. Lastly, when using the predetermined design set, the Vandermonde matrix $\mathsf{M}(\Phi, \mathcal{Y})$ in Definition 2.1 exhibits a unique structure that reduces the matrix inversion cost to $\mathcal{O}(d^2)$ from $\mathcal{O}(d^3)$.

In addition to the coordinate-basis selection of the points, ASTRO-DF-C also skips the criticality step in the model construction (Shashaani *et al.*, 2018). The criticality step is a loop within each iteration that enforces additional shrinkage of the trust region until $\|\nabla M_k(\boldsymbol{X}_k)\|$ is smaller than a factor of the trust-region radius Δ_k. Every time the trust-region radius is shrunk due to this criticality check, more points may be needed in the contracted region to fit a new local model, which hurts the work complexity in each iteration. Instead, ASTRO-DF-C only checks the criteria $\mu \|\nabla M_k(\boldsymbol{X}_k)\| \geq \Delta_k$ once at the time of updating the trust region for the next iteration to determine whether to accept the model candidate $\tilde{\boldsymbol{X}}_{k+1}$ if the interpolation candidate $\hat{\boldsymbol{X}}_{k+1}$ fails to pass its sufficient reduction test.

3.2. Updating the next incumbent

The next incumbent is determined upon estimating the function at the subproblem's minimizer. If the best among

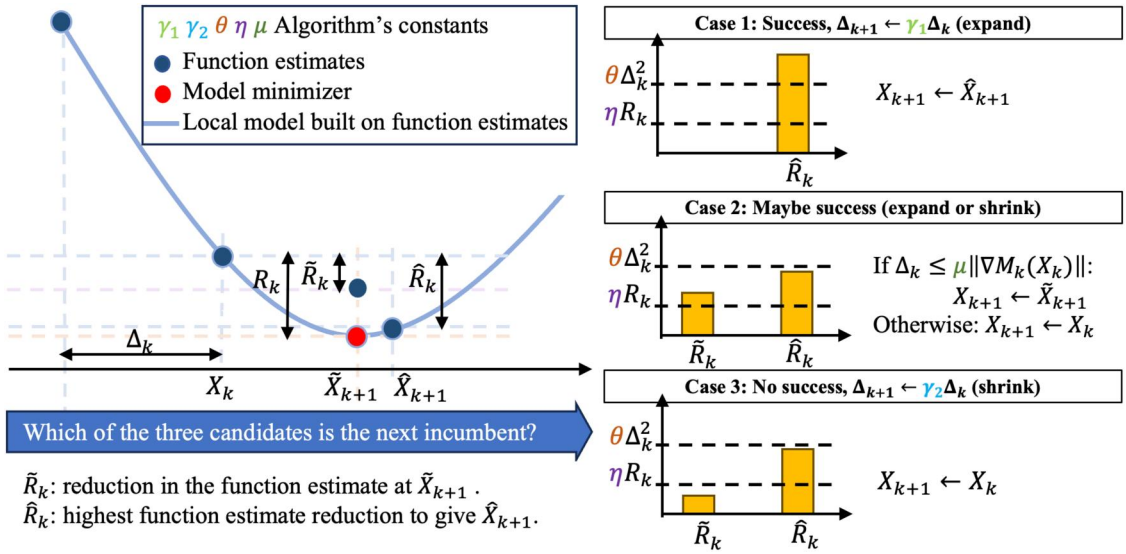

Figure 3. Updating of the incumbent is through one of three possible cases. ASTRO-DF only has Case 2 and Case 3. However, the direct search provides a new possibility for success by using the existing function estimates that are lower than that of the model minimizer, provided they satisfy a minimum reduction of $\theta\Delta_k^2$. The functionality of the algorithmic constants in this process and in updating the next trust region is also illustrated.

$2d + 2$ points ($2d$ interpolation points, one incumbent point, and one model candidate point) is one of the $2d$ interpolation points that satisfies a reduction of at least $\theta\Delta_k^2$ in the estimated function value, that point becomes the next incumbent; otherwise, the model candidate is accepted if the decrease in the estimated function value is at least η times the reduction in the model and the model satisfies $\mu\|\nabla M_k(\boldsymbol{X}_k)\| \geq \Delta_k$. Figure 3 visualizes the above procedure.

Bypassing model construction and choosing the incumbent from the coordinate-basis design points *directly* turns this process into a direct search, somewhat similar to algorithms such as NOMAD (Audet *et al.*, 2021). Direct search is a numerical optimization procedure that, unlike trust region or line search methods that use derivative information, only relies on function evaluations of a design set to determine whether to move to a new incumbent. In ASTRO-DF, choosing a design point that outperforms the model candidate point, albeit with a threshold of sufficient reduction, as the next incumbent will increase *the probability of success* (see Theorem 4.7). The threshold of sufficient reduction is necessary because a model candidate point is considered qualified if the model improves the function estimate by a factor of Δ_k^2, provided that the model is certified. In the absence of a model, one needs an alternative criterion to maintain the quality of the next incumbent. However, interestingly, the inclusion of direct search makes it a possibility that we accept a point with a reduction less than ηR_k if $\eta R_k > \theta\Delta_k^2$. Meanwhile, ASTRO-DF would reject such a point. This renders acceptance criteria for the new incumbent less strict. The resulting improved probability of success takes effect in the early iterations, hence helping the finite-time performance. As the iterations progress, the model tends to approximate the function more accurately in smaller neighborhoods, which reduces the likelihood of the model candidate point underperforming the interpolation points.

3.3. Simulation budget

The deterministic lower bound sequence for the sample sizes, i.e., $\{\lambda_k, k \in \mathbb{N}\}$ now grows logarithmically instead of linearly with k. This leads to a slower growth in the minimum sample size and savings of the simulation budget throughout. Other than the slower rate of decrease in $\lambda_k T$ the stopping time sample size (6) is identical to the original version of ASTRO-DF, continuously adding an independent new simulation run at $\boldsymbol{X}_k$ until the standard error falls below the slightly deflated (with $\sqrt{\lambda_k}$) fourth power of the trust-region radius.

4. Convergence and complexity analysis

This section presents the convergence and complexity analysis of ASTRO-DF-C. We first list the assumptions and useful results concerning ASTRO-DF. The needed assumptions to obtain the convergence and complexity results are standard and not restrictive, allowing for the use of the Bernstein inequality to bound the tail of sums of subexponential random variables that are not independent of one another.

Assumption 1. *(function) The function f is twice continuously differentiable in an open domain Ω, ∇f is Lipschitz continuous in Ω with constant $\kappa_{Lg} > 0$, and $\nabla^2 f$ is Lipschitz continuous in Ω with constant $\kappa_L > 0$.*

For the sake of clarity in notation, we henceforth replace $E(\boldsymbol{X}_k^i, \xi_j)$ with $E_{k,j}^i$, $N(\boldsymbol{X}_k^i)$ with N_k^i, and $\bar{E}(\boldsymbol{X}_k^i, N_k^i)$ with $\bar{E}_k^i$. For the center point $\boldsymbol{X}_k$, we interchangeably use $\boldsymbol{X}_k^0$.

Assumption 2. *(stochastic error) The Monte Carlo oracle generates iid random variables $F(\boldsymbol{X}_k^i, \xi_j) = f(\boldsymbol{X}_k^i) + E_{k,j}^i$ with $E_{k,j}^i \in \mathcal{F}_{k,j}$, where $\mathcal{F}_k := \mathcal{F}_{k,0} \subset \mathcal{F}_{k,1} \subset \ldots \subset \mathcal{F}_{k+1}$ for all k. The design set is $\{\boldsymbol{X}_k^i\}_{i=0,1,\ldots,p} \in \mathcal{F}_{k-1}$. Then the stochastic errors $E_{k,j}^i$ are independent of $\mathcal{F}_{k-1}$, $\mathbb{E}[E_{k,j}^i|\mathcal{F}_{k,j-1}] = 0$, and*

there exists $\sigma^2 > 0$ and $b > 0$ such that for a fixed n,

$$\frac{1}{n}\sum_{j=1}^{n}\mathbb{E}\left[|E_{k,j}^i|^m|\mathcal{F}_{k,j-1}\right] \leq \frac{m!}{2}b^{m-2}\sigma^2, \forall m = 2,3,\ldots, \forall k.$$

The above is known as the Bernstein condition. Furthermore, for a given incumbent at iteration k, there exists a large $c_0 > 0$ such that for all $\lambda_k \leq n \leq N_k^i$,

$$\limsup_{c\to\infty}\sup_{c_0\leq t\leq c}\frac{\mathbb{P}\left\{\sum_{j=1}^{n-1}|E_{k,j}^i| > c-t||E_{k,n}^i| = t\right\}}{\mathbb{P}\left\{\sum_{j=1}^{n-1}|E_{k,j}^i| > c-t\right\}} < \infty. \quad (7)$$

In other words, the ratio above is $\mathcal{O}(1)$ uniformly for all $t \in [c_0, c]$.

The criteria in (7) intuitively mean we cannot have extreme positive dependence between the latest observed error's magnitude and the previous total observed error magnitudes. That is, the probability of the event $\left\{(\sum_{j=1}^{n-1}|E_{k,j}^i| > c-t)\cap(|E_{k,n}^i| \in \mathcal{B}(t))\right\}$ cannot be much larger than $\mathbb{P}\{\sum_{j=1}^{n-1}|E_{k,j}^i| > c-t\}\mathbb{P}\{|E_{k,n}^i| \in \mathcal{B}(t)\}$, where $\mathbb{P}\{|E_{k,n}^i| \in \mathcal{B}(t)\}$ is to be understood as the probability of $|E_{k,n}^i|$ lying in some open interval containing t, for any $t < c$ larger than c_0. For correctness, when $\mathbb{P}\{|E_{k,n}^i| \in \mathcal{B}(t)\}$ is not possible, the conditional probability is simply understood as zero. Ko and Tang (2008) show that this assumption allows a wide range of dependence structures. Importantly, this assumption enables characterization of the tail probability of a sequence of dependent random variables that each have a sub-exponential tail behavior for the Bernstein inequality bounds to be applicable. The criteria for the stochastic error in Assumption 2 are less stringent than boundedness or sub-Gaussian behavior. See (Ha *et al.*, 2023) for more details.

Assumption 3. *(model) There exists some constant $\kappa_{fcd} \in (0,1]$ such that for all k,*

$$M_k(\boldsymbol{X}_k) - M_k(\boldsymbol{X}_k + \boldsymbol{S}_k) \geq \kappa_{fcd}[M_k(\boldsymbol{X}_k) - M_k(\boldsymbol{X}_k + \boldsymbol{S}_k^c)],$$

where $\boldsymbol{S}_k^c$ is the Cauchy step. Additionally, there exists $\kappa_{\mathsf{H}} \in (0,\infty)$ such that $\|\mathsf{H}_k\| \leq \kappa_{\mathsf{H}}$ for all k with probability 1.

Assumption 4. *(sample size) The "lower-bound sequence" $\{\lambda_k\}$ on the adaptive sample sizes satisfies $\lambda_k = \mathcal{O}((\log k)^{1+\epsilon_\lambda})$ for some $\epsilon_\lambda \in (0,1)$.*

4.1. Useful existing results

In this section, we introduce several useful results to obtain the almost sure convergence and complexity of ASTRO-DF with coordinate direct search. The first result implies that the estimation error with adaptive sampling is sufficiently small eventually (for large enough iterations) with probability 1.

Theorem 4.1. *(Theorem 2 by Ha et al. (2023)) Let Assumptions 2 and 4 hold. Then for any $c > 0$ and sufficiently large k, $|\bar{E}_k^i| \leq c\Delta_k^2$ almost surely. In other words,*

$$\mathbb{P}\left\{|\bar{E}_k^i| \geq c\Delta_k^2 \ i.o.\right\} = 0.$$

For a sketch of this proof, observe that the subexponential behavior of the stochastic error specified in Assumption 2 makes it possible to apply the Bernstein inequality for the average of stochastic error absolute values. Thus,

$$\mathbb{P}\left\{\frac{1}{N_k}\sum_{j=1}^{N_k}|E_{k,j}^i| > c|\mathcal{F}_{k-1}\right\} \leq \sum_{n\geq\lambda_k}\exp\left(-n\frac{c^2}{2(cb+\sigma^2)}\right),$$

since $N_k \geq \frac{\sigma^2\lambda_k}{2\kappa^2\Delta_k^2}$ from (6) for large enough k. The logarithmic λ_k here then renders this probability summable in k. Invoking the Borel–Cantelli Lemma then completes the proof.

Remark 1. Bernstein's inequality implies that $\mathbb{P}\{|\bar{E}_k^i| \geq c\Delta_k^2|\mathcal{F}_k\} \leq \alpha k^{-1+\epsilon_\lambda}$ for any $c > 0$ and some $\alpha > 0$; i.e., the estimate is sufficiently accurate with a probability that tends to one as $k \to \infty$. In contrast to the assumption of probabilisitcally accurate estimates with a fixed probability and fixed accuracy threshold (Chen *et al.*, 2018), the estimators' property here is ensured as a result of the adaptive sampling rule.

The second existing result characterizes the stochastic model error with estimation error when the stochastic model is constructed by Definition 2.2.

Lemma 4.2. *(Lemma 2.9 by Shashaani et al. (2018) and Proposition 3.1 by Ragonneau and Zhang (2023)) Let Assumption 1 hold and let $M_k(\cdot)$ be a stochastic quadratic interpolation model with diagonal Hessian of f on $\mathcal{B}(\boldsymbol{X}_k;\Delta_k)$. Let $m_k(\cdot)$ be the corresponding deterministic polynomial interpolation model of f on $\mathcal{B}(\boldsymbol{X}_k;\Delta_k)$. Then for all $\boldsymbol{x} \in \mathcal{B}(\boldsymbol{X}_k;\Delta_k)$,*

$$|M_k(\boldsymbol{x}) - m_k(\boldsymbol{x})| \leq (2d+1)\max_{i=0,1,\ldots,2d}|\bar{F}(\boldsymbol{X}_k^i, n(\boldsymbol{X}_k^i)) - f(\boldsymbol{X}_k^i)|.$$

Remark 2. A direct observation from Lemma 4.2, Theorem 4.4, and Appendix A is that

$$\mathbb{P}\{\|\nabla f(\boldsymbol{x}) - \nabla M_k(\boldsymbol{x})\| \geq (c_{eg}+\kappa_{eg})\Delta_k \text{ and}$$
$$|f(\boldsymbol{x}) - M_k(\boldsymbol{x})| \geq (c_{ef}+\kappa_{ef})\Delta_k^2|\mathcal{F}_k\} \leq \alpha' k^{-1+\epsilon_\lambda},$$

for some $\alpha' > 0$, $c_{eg} > 0$, and $c_{ef} > 0$ and for all $\boldsymbol{x} \in \mathcal{B}(\boldsymbol{X}_k;\Delta_k)$. In other words, the stochastic model is probabilistically fully linear with a probability that goes to one as $k \to \infty$. Again, despite the resemblance to the assumption of probabilistically fully linear models with fixed probability (Chen *et al.*, 2018), this result is ensured without overburdening the computation, as we will see in the obtained complexity results.

4.2. Model quality

To prove consistency, we first look at the model quality with the suggested design set in Definition 2.2. In general, the stochastic polynomial interpolation model satisfies that the model gradient norm is eventually $\mathcal{O}(\Delta_k)$-accurate for any point within the trust region of size Δ_k with probability 1 (Shashaani *et al.*, 2018). The following theorem gives a

bound on the stochastic model error that holds for any point within the trust region.

Theorem 4.3. *Let Assumptions 1 and 2 hold, and $M_k(\boldsymbol{x})$ be a stochastic quadratic model of f on $\mathcal{B}(\boldsymbol{X}_k;\Delta_k)$ with diagonal Hessian built on $\bar{F}(\boldsymbol{X}_k^i, n(\boldsymbol{X}_k^i)) = f(\boldsymbol{X}_k^i) + \bar{E}_k^i$, for $i = 0, 1, ..., 2d$ on iteration k. With $\nabla M_k(\boldsymbol{x}) = (\boldsymbol{x} - \boldsymbol{X}_k)^\top \mathsf{H}_k + \boldsymbol{G}_k$ and $\kappa_{eg1} = \frac{5\sqrt{2d}}{2}(\kappa_{Lg} + \kappa_{\mathsf{H}})$, we can uniformly bound the model gradient error for all points $\boldsymbol{x}$ in $\mathcal{B}(\boldsymbol{X}_k;\Delta_k)$ by*

$$\|\nabla M_k(\boldsymbol{x}) - \nabla f(\boldsymbol{x})\| \le \kappa_{eg1}\Delta_k + \frac{\sqrt{\sum_{i=1}^{2d}\left(\bar{E}_k^i - \bar{E}_k^0\right)^2}}{\Delta_k}.$$

Proof. With a stochastic quadratic models with diagonal Hessian, the theorem holds following the same steps of Appendix B by Shashaani *et al.* (2018). ■

We next investigate the model gradient error at the current iterate as a special case. At the center point, we can get a tighter upper bound, as with a central finite difference. Theorem 4.4 shows the deterministic part of the model gradient error at the incumbent is $\mathcal{O}(\Delta_k^2)$.

Theorem 4.4. *Let Assumptions 1 and 2 hold, and the interpolation model $M_k(\boldsymbol{x})$ of f be a stochastic quadratic model with diagonal Hessian constructed using $\bar{F}(\boldsymbol{X}_k^i, n(\boldsymbol{X}_k^i)) = f(\boldsymbol{X}_k^i) + \bar{E}_k^i$, for $i = 0, 1, ..., 2d$. Then, with $\kappa_{eg2} = \frac{\sqrt{d}}{6}\kappa_L$, we can uniformly bound the model gradient error at the center point by*

$$\|\boldsymbol{G}_k - \nabla f(\boldsymbol{X}_k)\| \le \kappa_{eg2}\Delta_k^2 + \frac{\sqrt{\sum_{i=1}^{d}\left(\bar{E}_k^i - \bar{E}_k^{i+d}\right)^2}}{2\Delta_k}.$$

Proof. We begin by decomposing $\boldsymbol{G}_k = \nabla f(\boldsymbol{X}_k) + E^g(\boldsymbol{X}_k) + e^g(\boldsymbol{X}_k)$, where $\boldsymbol{E}_k^g := \boldsymbol{G}_k - \boldsymbol{g}_k$ and $\boldsymbol{e}_k^g := \boldsymbol{g}_k - \nabla f(\boldsymbol{X}_k)$ are the stochastic model error and deterministic model error, respectively, and $\boldsymbol{g}_k$ is the deterministic model gradient. Then we can obtain $\boldsymbol{G}_k$ by finding $\boldsymbol{\beta}$ such that $\tilde{\mathsf{M}}(\Phi, \mathcal{X}_k)\boldsymbol{\beta} = \tilde{\boldsymbol{F}}$, where

$$\tilde{\mathsf{M}}(\Phi, \mathcal{X}_k) = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ 1 & \Delta_k & 0 & \cdots & 0 & \Delta_k^2/2 & 0 & \cdots & 0 \\ 1 & 0 & \Delta_k & \cdots & 0 & 0 & \Delta_k^2/2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & \cdots & \Delta_k & 0 & 0 & \cdots & \Delta_k^2/2 \\ 1 & -\Delta_k & 0 & \cdots & 0 & \Delta_k^2/2 & 0 & \cdots & 0 \\ 1 & 0 & -\Delta_k & \cdots & 0 & 0 & \Delta_k^2/2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & \cdots & -\Delta_k & 0 & 0 & \cdots & \Delta_k^2/2 \end{bmatrix};$$

$$\tilde{\boldsymbol{F}} = \begin{bmatrix} \bar{F}(\boldsymbol{X}_k^0, N_k^0) \\ \bar{F}(\boldsymbol{X}_k^1, N_k^1) \\ \vdots \\ \bar{F}(\boldsymbol{X}_k^{2d}, N_k^{2d}) \end{bmatrix}.$$

As a result, the ith element of the gradient estimate is obtained by

$$[\boldsymbol{G}_k]_i = \frac{f(\boldsymbol{X}_k + \boldsymbol{e}_i\Delta_k) - f(\boldsymbol{X}_k - \boldsymbol{e}_i\Delta_k)}{2\Delta_k} + \frac{\bar{E}_k^i - \bar{E}_k^{i+d}}{2\Delta_k}. \tag{8}$$

By Taylor's theorem,

$$\|\boldsymbol{g}_k - \nabla f(\boldsymbol{X}_k)\| \le \frac{\sqrt{d}}{6}\kappa_L\Delta_k^2,$$

with (8) like the central finite difference (Berahas Cao, Choromanski and Scheinberg, 2021; Coope and Tappenden, 2021). Hence, we obtain $[\boldsymbol{E}_k^g]_i = (\bar{E}_k^i - \bar{E}_k^{i+d})/2\Delta_k$ for any $i \in \{1, ..., d\}$, which implies that $\|\boldsymbol{E}_k^g\| = \sqrt{\sum_{i=1}^{d}(\bar{E}_k^i - \bar{E}_k^{i+d})^2}/(2\Delta_k)$. As a result,

$$\begin{aligned}\|\boldsymbol{G}_k - \nabla f(\boldsymbol{X}_k)\| &\le \|\boldsymbol{g}_k - \nabla f(\boldsymbol{X}_k)\| + \|\boldsymbol{G}_k - \boldsymbol{g}_k\| \\ &\le \frac{\sqrt{d}}{6}\kappa_L\Delta_k^2 + \frac{\sqrt{\sum_{i=1}^{d}(\bar{E}_k^i - \bar{E}_k^{i+d})^2}}{2\Delta_k}.\end{aligned}$$

■

4.3. Convergence

For the remainder of this article, we define the set $\mathcal{K} = \{k \in \mathbb{N} : \text{iteration } k \text{ is successful}\}$.

Theorem 4.5. *Let Assumptions 1-4 hold. Then, $\Delta_k \overset{w.p.1}{\to} 0$ as $k \to \infty$.*

Proof. If $\mathcal{K}$ is finite, the trust-region radius tends to zero due to shrinkage over infinitely many unsuccessful iterations, making the statement of the theorem hold trivially. Thus, we consider that $\mathcal{K}$ is infinite. The candidate incumbent can be one of the points among the design sets or the trust-region subproblem minimizer. Let us define two different sets $\mathcal{K}_{ds} := \mathcal{K} \cap \{k \in \mathbb{N} : \boldsymbol{X}_{k+1} = \hat{\boldsymbol{X}}_{k+1}\}$ and $\mathcal{K}_{tr} := \mathcal{K} \cap \{k \in \mathbb{N} : \boldsymbol{X}_{k+1} = \tilde{\boldsymbol{X}}_{k+1}\}$. Then we have, for any $k \in \mathcal{K}_{ds}$,

$$\hat{R}_k = \bar{F}(\boldsymbol{X}_k, N_k) - \bar{F}(\hat{\boldsymbol{X}}_{k+1}, \hat{N}_{k+1}) = M_k(\boldsymbol{X}_k) - M_k(\hat{\boldsymbol{X}}_{k+1}) \ge \theta\Delta_k^2,$$

and for any $k \in \mathcal{K}_{tr}$,

$$\begin{aligned}\tilde{R}_k = \bar{F}(\boldsymbol{X}_k, N_k) - \bar{F}(\tilde{\boldsymbol{X}}_{k+1}, \tilde{N}_{k+1}) &\ge \eta(M_k(\boldsymbol{X}_k) - M_k(\tilde{\boldsymbol{X}}_{k+1})) \\ &\ge \frac{\eta\kappa_{fcd}}{2}\frac{\Delta_k}{\mu}\left(\frac{\Delta_k}{\mu\kappa_{\mathsf{H}}} \wedge \Delta_k\right) \\ &\ge \left(\frac{\eta\kappa_{fcd}}{2\mu}\left(\frac{1}{\mu\kappa_{\mathsf{H}}} \wedge 1\right)\right)\Delta_k^2.\end{aligned}$$

Then, for any $k \in \mathcal{K}$,

$$\begin{aligned}\theta' \sum_{k\in\mathcal{K}} \Delta_k^2 &\le \sum_{k\in\mathcal{K}} (f(\boldsymbol{X}_k) - f(\boldsymbol{X}_{k+1}) + \bar{E}_k - \bar{E}_{k+1}) \\ &\le f(\boldsymbol{x}_0) - f^* + \sum_{k=0}^{\infty}(|\bar{E}_k| + |\bar{E}_{k+1}|),\end{aligned}$$

where $\theta' = \left(\frac{\eta\kappa_{fcd}}{2\mu}\left(\frac{1}{\mu\kappa_H}\wedge 1\right)\wedge\theta\right)$. Let $\mathcal{K} = \{k_1, k_2, \ldots\}$, $k_0 = -1$, and $\Delta_{-1} = \Delta_0/\gamma_2$. Then from the fact that $\Delta_k \leq \gamma_1\gamma_2^{k-k_i-1}\Delta_{k_i}$ for $k = k_i+1, \ldots, k_{i+1}$ and each i, we obtain

$$\sum_{k=k_i+1}^{k_{i+1}}\Delta_k^2 \leq \gamma_1^2\Delta_{k_i}^2\sum_{k=k_i+1}^{k_{i+1}}\gamma_2^{2(k-k_i-1)} \leq \gamma_1^2\Delta_{k_i}^2\sum_{k=0}^{\infty}\gamma_2^{2k} = \frac{\gamma_1^2}{1-\gamma_2^2}\Delta_{k_i}^2.$$

Then, we have

$$\sum_{k=0}^{\infty}\Delta_k^2 \leq \frac{\gamma_1^2}{1-\gamma_2^2}\sum_{i=0}^{\infty}\Delta_{k_i}^2 < \frac{\gamma_1^2}{1-\gamma_2^2}\left(\frac{\Delta_0^2}{\gamma_2^2}+\frac{f(\boldsymbol{x}_0)-f^*+E'_{0,\infty}}{\theta'}\right),$$

where $E'_{i,j} = \sum_{k=i}^{j}(|\bar{E}_k| + |\bar{E}_{k+1}|)$. By Theorem 4.1, there must exist a sufficiently large K_Δ such that $|\bar{E}_k| + |\bar{E}_{k+1}| < c_\Delta\Delta_k^2$ for any given $c_\Delta > 0$ and every $k \geq K_\Delta$. Then we obtain

$$\sum_{k=0}^{K_\Delta-1}\Delta_k^2 + \sum_{k=K_\Delta}^{\infty}\Delta_k^2 < \frac{\gamma_1^2}{1-\gamma_2^2}\left(\frac{\Delta_0^2}{\gamma_2^2}+\frac{f(\boldsymbol{x}_0)-f^*+E'_{0,K_\Delta-1}+E'_{K_\Delta,\infty}}{\theta'}\right),$$

and $E'_{K_\Delta,\infty} = \sum_{k=K_\Delta}^{\infty}|\bar{E}_k| + |\bar{E}_{k+1}| < \sum_{k=K_\Delta}^{\infty}c_\Delta\Delta_k^2$. As a result, we get

$$\sum_{k=K_\Delta}^{\infty}\Delta_k^2 < \frac{\gamma_1^2}{1-\gamma_2^2}\left(\frac{\Delta_0^2}{\gamma_2^2}+\frac{f(\boldsymbol{x}_0)-f^*+E'_{0,K_\Delta-1}}{\theta'}\right)\left(1-\frac{\gamma_1^2}{1-\gamma_2^2}\frac{c_\Delta}{\theta'}\right)^{-1} < \infty. \tag{9}$$

Therefore, $\Delta_k \to^{wp1} 0$ as $k \to \infty$. ■

Now we prove the almost sure convergence of ASTRO-DF-C.

Theorem 4.6. *Let Assumptions 1-4 hold. Then, $\|\nabla f(\boldsymbol{X}_k)\| \xrightarrow{w.p.1} 0$ as $k \to \infty$.*

Proof. Let us define the set

$$\mathcal{V} := \left\{\exists \text{ a subsequence } \{k_j\} \right.$$
$$\left. \text{s.t. } \left(\frac{\Delta_{k_j}}{\|\boldsymbol{G}_{k_j}\|} \leq \left(\frac{\eta'}{2\kappa_{ef}}\wedge\mu\right)\right)\cap(\hat{\rho}_k < \eta)\right\},$$

where $\eta' = 6^{-1}(1-\eta)\kappa_{fcd}((\mu\kappa_H)^{-1}\wedge 1)$. Then, we have $\mathbb{P}\{\mathcal{V}\} = 0$. The proof trivially follows from Lemma 5.2 in Shashaani *et al.* (2018) by considering that $\{\mu\|\boldsymbol{G}_{k_j}\| \geq \Delta_{k_j}\}$ is now in the set $\mathcal{V}$, which was ensured by the criticality step by Shashaani *et al.* (2018). We also have $\|\boldsymbol{G}_k - \nabla f(\boldsymbol{X}_k)\| \xrightarrow{wp1} 0$ as $k \to \infty$. The proof follows from that of Lemma 5.4 by Shashaani *et al.* (2018), considering that we always use Δ_k for iteration k. Now we will prove that Algorithm 1 obtains $\liminf\|\nabla f(\boldsymbol{X}_k)\| \xrightarrow{wp1} 0$ as $k \to \infty$ with $\|\boldsymbol{G}_k - \nabla f(\boldsymbol{X}_k)\| \xrightarrow{wp1} 0$ as $k \to \infty$ and Theorem 4.5. For the purpose of arriving at a contradiction, suppose the set

$$\mathcal{D}_g = \{\omega : \exists\kappa_{lbg}(\omega), k_g(\omega) > 0 \text{ s.t. } \|\boldsymbol{G}_k\| \geq \kappa_{lbg}(\omega)\,\forall\, k > k_g(\omega)\}$$

has positive measure. Due to the assumptions of the theorem, we can find a set $\mathcal{D}_d$ of sample paths such that $\mathbb{P}\{\mathcal{D}_d\} = 1$, and such that for each $\omega \in \mathcal{D}_d$, $\Delta_k(\omega) \to 0$ and $\omega \in \mathcal{V}^c$. Let $\omega \in \mathcal{D}_g \cap \mathcal{D}_d$. Then either $\|\boldsymbol{G}_k(\omega)\| < (\eta'(2\kappa_{ef})^{-1}\wedge\mu)^{-1}\,\Delta_k(\omega)$ or $\hat{\rho}_k(\omega) > \eta$ for large enough k. Since Δ_k goes to zero almost surely, $\|\boldsymbol{G}_k(\omega)\| < (\eta'(2\kappa_{ef})^{-1}\wedge\mu)^{-1}\Delta_k(\omega)$ cannot be true for large enough k. Therefore, for any $\omega \in \mathcal{D}_g\cap\mathcal{D}_d$, it must be true that $\hat{\rho}_k(\omega) \geq \eta$ for large enough k. In other words, the iterations in sample path $\omega \in \mathcal{D}_g\cap\mathcal{D}_d$ are eventually successful.

Now let $K_s(\omega) > 0$ be such that $K_s(\omega) - 1$ is the last unsuccessful iteration in sample-path $\omega \in \mathcal{D}_g \cap \mathcal{D}_d$, that is, k is a successful iteration if $k \geq K_s(\omega)$. Next, $\Delta_k \geq \Delta_{K_g(\omega)\vee K_s(\omega)}$ contradicts the observation $\Delta_k(\omega) \to 0$. We conclude that $\mathbb{P}\{\mathcal{D}_g\} = 0$ and that $\liminf_{k\to\infty}\|\boldsymbol{G}_k\| = 0$ almost surely. This along with the fact that $\|\boldsymbol{G}_k - \nabla f(\boldsymbol{X}_k)\| \xrightarrow{wp1} 0$ as $k \to \infty$, implies $\liminf_{k\to\infty}\|\nabla f(\boldsymbol{X}_k)\| = 0$ almost surely. Then, the almost sure convergence of ASTRO-DF-C follows from $\liminf_{k\to\infty}\|\boldsymbol{G}_k\| = 0$ almost surely, and Theorem 4.5. The proof is completed by trivially following steps in Theorem 5.5 in Shashaani *et al.* (2018) and considering that $\hat{R}_k \geq \theta\Delta_k^2$. ■

Now we prove that ASTRO-DF-C has a higher probability of success compared with ASTRO-DF. Although this result is not utilized in our current complexity results, it demonstrates that the progress achieved by integrating a coordinate direct search is at least as good as the progress made without that feature, thereby supporting its positive impact on the finite-time performance of the algorithm.

Theorem 4.7. *Let Assumptions 1-4 hold. Then for a given incumbent $\boldsymbol{X}_k \in \mathbb{R}^d$ in iteration k, having a successful iteration with ASTRO-DF-C is at least as probable as having a successful iteration with ASTRO-DF.*

Proof. Let us define two events

$$\mathcal{R}_k := \{\omega \in \Omega : k(\omega) \text{ is successful with } \hat{\boldsymbol{X}}_{k+1}(\omega)|\boldsymbol{X}_k(\omega) = \boldsymbol{x}_k\},$$
$$\mathcal{S}_k := \{\omega \in \Omega : k(\omega) \text{ is successful with } \tilde{\boldsymbol{X}}_{k+1}(\omega)|\boldsymbol{X}_k(\omega) = \boldsymbol{x}_k\}.$$

Then $p_{k,s} := \mathbb{P}\{\mathcal{S}_k\} = \mathbb{P}\{\mathcal{S}_k|\mathcal{R}_k\}\mathbb{P}\{\mathcal{R}_k\} + \mathbb{P}\{\mathcal{S}_k|\mathcal{R}_k^\complement\}\mathbb{P}\{\mathcal{R}_k^\complement\}$, is the probability of having successful iteration k with ASTRO-DF. In contrast, ASTRO-DF-C sequentially compares $\hat{\boldsymbol{X}}_{k+1}(\omega)$ and $\tilde{\boldsymbol{X}}_{k+1}(\omega)$. If the iteration k is unsuccessful with $\hat{\boldsymbol{X}}_{k+1}(\omega)$, $\tilde{\boldsymbol{X}}_{k+1}(\omega)$ will still be considered as the next incumbent, i.e., $p_{k,r} = \mathbb{P}\{\mathcal{R}_k\} + \mathbb{P}\{\mathcal{S}_k|\mathcal{R}_k^\complement\}$, where $p_{k,r}$ is the probability of having successful iteration k with ASTRO-DF-C. As a result, $p_{k,r} \geq p_{k,s}$. ■

Remark 3. Since the sequence $\{\boldsymbol{X}_k\}$ is dependent on previous steps, we cannot directly compare ASTRO-DF with ASTRO-DF-C using Theorem 4.7. Nonetheless, Theorem 4.7 implies that, given an incumbent $\boldsymbol{X}_k$, ASTRO-DF-C has a higher likelihood of achieving success, preventing Δ_k from

becoming too small too quickly. This enables the algorithm to save significant budget due to $N_k = O(\Delta_k^{-4})$.

4.4. Complexity

While the iteration complexity and work complexity of ASTRO-DF has been extensively studied by Ha *et al.* (2023), our focus in this section is on the refinements and their impact on the complexity analysis. Specifically, we will examine how these refinements affect the algorithm's computational efficiency.

The following lemma proves that when Δ_k is too small relative to $\|\nabla f(\boldsymbol{X}_k)\|$, the iteration k becomes successful almost surely for sufficiently large k.

Lemma 4.8. *Let Assumptions 1-4 hold and $\epsilon > 0$ be given. Then there exists a constant $c_{lb} > 0$ such that with probability 1*

$$(\Delta_k < c_{lb}\epsilon) \cap (\|\nabla f(\boldsymbol{X}_k)\| > \epsilon) \text{ for sufficiently large } k \Rightarrow k \in \mathcal{K}.$$

Proof. Let the solver sample path ω be fixed. Define a constant c_E as

$$c_E = \frac{1}{2d+2}\left(\frac{\kappa_{fcd}(1-\eta)}{2\mu}\left(\frac{1}{\mu\kappa_{\mathsf{H}}} \wedge 1\right) - \kappa_{ef}\right). \tag{10}$$

We can find $K_E(\omega) > 0$ such that $k \geq K_E(\omega)$ implies $|\bar{E}_k(\omega)| \leq c_E\Delta_k^2(\omega)$, which also holds for all design points visited during iteration k by Theorem 4.1. Then Theorem 4.4 states for $k \geq K_E(\omega)$,

$$\|\boldsymbol{G}_k(\omega) - \nabla f(\boldsymbol{X}_k(\omega))\| \leq \frac{\sqrt{d}}{6}\kappa_L\Delta_k^2(\omega) + \frac{\sqrt{d}}{\sqrt{2}}c_E\Delta_k(\omega).$$

Let c_{lb} be such that

$$\frac{1}{c_{lb}} > \frac{1}{\mu} + \frac{\sqrt{d}}{6}\kappa_L\Delta_{\max} + \frac{\sqrt{d}}{\sqrt{2}}c_E.$$

Then, if $\Delta_k(\omega) < c_{lb}\epsilon$, we get

$$\begin{aligned}\|\boldsymbol{G}_k(\omega)\| &\geq \|\nabla f(\boldsymbol{X}_k(\omega))\| - \|\boldsymbol{G}_k(\omega) - \nabla f(\boldsymbol{X}_k(\omega))\| \\ &> \left(\frac{1}{c_{lb}} - \frac{\sqrt{d}}{6}\kappa_L\Delta_{\max} + \frac{\sqrt{d}}{\sqrt{2}}c_E\right)\Delta_k(\omega) > \frac{1}{\mu}\Delta_k(\omega),\end{aligned}$$

where we have used $\|\nabla f(\boldsymbol{X}_k(\omega))\| > \epsilon$.

This result confirms that the model quality is eventually good whenever the trust region becomes sufficiently small. To complete the proof we need to show the model will lead to success; a sufficient condition for that is $\tilde{R}_k(\omega) > \eta R_k(\omega)$. For ease of exposition we drop ω in the final step of the proof:

$$\begin{aligned}\left|1 - \frac{\tilde{R}_k}{R_k}\right| &= \left|\frac{\bar{F}(\tilde{\boldsymbol{X}}_{k+1}, \tilde{N}_{k+1}) - M_k(\tilde{\boldsymbol{X}}_{k+1})}{M_k(\boldsymbol{X}_k) - M_k(\tilde{\boldsymbol{X}}_{k+1})}\right| \\ &\leq \frac{|\bar{E}(\tilde{\boldsymbol{X}}_{k+1})| + |f(\tilde{\boldsymbol{X}}_{k+1}) - m_k(\tilde{\boldsymbol{X}}_{k+1})| + |m_k(\tilde{\boldsymbol{X}}_{k+1}) - M_k(\tilde{\boldsymbol{X}}_{k+1})|}{\frac{\kappa_{fcd}}{2\mu}\left(\frac{1}{\mu\kappa_{\mathsf{H}}} \wedge 1\right)\Delta_k^2} \\ &\leq \frac{c_E + \kappa_{ef} + (2d+1)c_E}{\frac{\kappa_{fcd}}{2\mu}\left(\frac{1}{\mu\kappa_{\mathsf{H}}} \wedge 1\right)} = 1 - \eta.\end{aligned} \tag{11}$$

■

Remark 4. In the proof of Lemma 4.8, the trust region lower bound constant c_{lb} depends on the coefficient of the tolerable estimation error c_E, specified in (10). The statement of the lemma conveys that if $\Delta_k(\omega) < c_{lb}\epsilon$ and $\|\nabla f(\boldsymbol{X}_k(\omega))\| > \epsilon$ for any given $\epsilon > 0$ and any $k > K_E(\omega)$, then the trust region will expand. So while ϵ can be chosen to be large such that Δ_k will not become too small to reach it, it does not cause any problem in the future complexity results. ASTRO-DF-C leads to a larger c_E since in (11) and (10), we now have 2d instead of $(d+1)(d+2)/2 - 1$ points. This means that compared to ASTRO-DF where $c_E = \mathcal{O}(d^{-2})$ and $c_{lb} = \mathcal{O}(d)$, this refined version will have $c_E = \mathcal{O}(d^{-1})$ and $c_{lb} = \mathcal{O}(\sqrt{d})$. Given that κ_{ef} and $\kappa_{\mathsf{H}} = \mathcal{O}(d\log d)$ are in our control (Appendix A and B), we can choose μ such that $c_E > 0$, i.e., $\mu^{-1} > \left(\sqrt{2\kappa_{\mathsf{H}}\kappa_{ef}} \wedge 2\kappa_{ef}\right)$. Nevertheless, the iteration complexity will be discussed for small enough ϵ.

Relying on Lemma 4.8, we now show the almost sure iteration complexity. This is stronger than the claim that the random variable $\epsilon^2 T_\epsilon$ is $\mathcal{O}_p(1)$.

Theorem 4.9. *Let Assumptions 1-4 hold. Then for sufficiently small $\epsilon > 0$, $c_T > 0$, and a positive integer valued random variable K,*

$$\mathbb{P}\{\epsilon^2 T_\epsilon \leq c_T + \epsilon^2 K\} = 1.$$

Proof. Let $f^* := \min_{\boldsymbol{x}\in\mathbb{R}^d} f(\boldsymbol{x}) > -\infty$ be the optimal function value, ω be fixed and $K_E(\omega)$ be the one defined in Lemma 4.8. We can find from Theorem 4.6 some $K_h(\omega)$ such that for all $k \geq K_h(\omega)$, we have that $f(\boldsymbol{X}_k(\omega)) < f_{st}^* + 1$, where f_{st}^* is the highest function value among the stationary points. Without loss of generality, given $\omega \in \Omega$, let $\epsilon_0(\omega) \in (0,1)$ be small enough such that except for a set of probability 0, the set $\mathcal{K}_{lb}(\omega) = \{K'_E(\omega) \leq k < T_\epsilon(\omega) : \Delta_k(\omega) < c_{lb}\epsilon\}$ is nonempty for all $\epsilon \leq \epsilon_0(\omega)$, where c_{lb} is defined in Lemma 4.8 and $K'_E(\omega) := (K_h(\omega) \vee K_E(\omega))$. This implies that $\Delta_k(\omega) \geq \gamma_2 c_{lb}\epsilon$ for all $K'_E(\omega) \leq k < T_\epsilon(\omega)$. For the remainder of the proof we will use the notation $f_k(\omega) := f(\boldsymbol{X}_k(\omega))$ for simplicity. Following the steps in the proof of Theorem 4.5 we have

$$\begin{aligned}\theta' \sum_{k\in\mathcal{K}} k \geq K'_E(\omega)\Delta_k^2(\omega) &\leq f_{K'_E(\omega)}(\omega) - f^* \\ &\quad + \sum_{k=K'_E(\omega)}^{\infty}(|\bar{E}_k(\omega)| + |\bar{E}_{k+1}(\omega)|),\end{aligned}$$

from which we obtain

$$\begin{aligned}\sum_{k=K'_E(\omega)}^{\infty}\Delta_k^2(\omega) &\leq \frac{\gamma_1^2}{1-\gamma_2^2}\sum_{i=0}^{\infty}\Delta_{k_i}^2(\omega) \\ &< \frac{\gamma_1^2}{1-\gamma_2^2}\left(\Delta_{\max^2} + \frac{f_{K'_E(\omega)}(\omega) - f^* + E'_{K'_E(\omega),\infty}(\omega)}{\theta'}\right),\end{aligned}$$

where $E'_{i,j}(\omega) = \sum_{k=i}^{j}(|\bar{E}_k(\omega)| + |\bar{E}_{k+1}(\omega)|)$ and k_i is the ith successful iteration after $K'_E(\omega)$. Then, by the definition of

$K'_E(\omega)$, we have $|\bar{E}_k(\omega)| \le c_E \Delta_k^2(\omega)$ for any $k \ge K'_E(\omega)$, which implies $E'_{K'_E(\omega),\infty}(\omega) \le 2c_E \sum_{k=K'_E(\omega)}^{\infty} \Delta_k^2(\omega)$. As a result, we obtain with small enough c_E,

$$\left(1 - \frac{2\gamma_1^2 c_E}{(1-\gamma_2^2)\theta'}\right) \sum_{k=K'_E(\omega)}^{\infty} \Delta_k^2(\omega) < \frac{\gamma_1^2}{1-\gamma_2^2}\left(\Delta_{\max^2} + \frac{f^*_{st} - f^* + 1}{\theta'}\right). \quad (12)$$

Now we can write

$$\sum_{k=K'_E(\omega)}^{\infty} \Delta_k^2(\omega) > \sum_{k=K'_E(\omega)}^{T_\epsilon(\omega)-1} \Delta_k^2(\omega) > \gamma_2^2 c_{lb}^2 \epsilon^2 (T_\epsilon(\omega) - K'_E(\omega)).$$

Then by (12), for all $\epsilon \le \epsilon_0(\omega)$, we have $\epsilon^2 T_\epsilon(\omega) < c_T + \epsilon^2 K'_E(\omega)$. ■

Note that Theorem 4.9 implies that $\limsup_{\epsilon \to 0} \epsilon^2 T_\epsilon(\omega) \le c_T$ for all $\omega \in \Omega$, that is, $\epsilon^2 T_\epsilon$ is bounded by a fixed value in the limit, almost surely.

Remark 5. A closer look in the proof of Theorem 4.9 suggests that ASTRO-DF-C has better iteration complexity than ASTRO-DF in the constant terms. This is because (dropping ω for ease of exposition) we have $\Delta_k \ge \gamma_2 c_{lb} \epsilon$ for all $K_E \le k < T_\epsilon$ with small enough ϵ. If for all $k < K_E$, there exists $C'_{lb} > 0$ (which is random and defined for sample path ω) such that $\Delta_k \ge \gamma_2 C'_{lb} \epsilon$, then $\sum_{k=0}^{T_\epsilon - 1} \Delta_k^2 > \gamma_2^2 c_{lb}^2 \epsilon^2 (T_\epsilon - K_E) + \gamma_2^2 C'^2_{lb} K_E \epsilon^2$. A larger value of C'_{lb} leads to a larger lower bound on the trust-region size, resulting in a smaller sample size (N_k) and a larger step size. This explains why the probability of having a successful iteration k is paramount. The two refinements play a crucial role in increasing c_{lb}, as stated in Remark 4 and Theorem 4.7. Additionally, under certain regularity conditions of random variable K_E such as finite first moment, Theorem 4.9 becomes equivalent to $\mathbb{E}[T_\epsilon] = \mathcal{O}(\epsilon^{-2})$; this is similar to the results proven in Blanchet *et al.* (2019). However, we attain this canonical rate without the assumptions of probabilistically-fully-linear models and independence, or using renewal theory (as used by Paquette and Scheinberg (2020) and Blanchet *et al.* (2019)).

5. Numerical results

In this section, we evaluate and compare simulation optimization solvers on problems from the SimOpt library (Eckman *et al.*, 2023).

The SimOpt library includes the MRG32k3a (L'Ecuyer *et al.*, 2002) pseudorandom-number generator and common random numbers for all solvers to manage uncertainties during search and evaluation and enable efficient comparisons. The SimOpt solver library includes various solvers like Nelder–Mead, Random Search, ALOE (Jin *et al.*, 2021), ADAM (Kingma and Ba, 2017), STRONG (Chang *et al.*, 2013), and STORM (Chen *et al.*, 2018). The SimOpt problem library consists of optimization problems where the simulation oracle provides the objective function value at specific points. Due to limited information on the objective function's structure, stochastic simulation oracles are preferred over deterministic problems with *added* stochastic error, as the latter leads to artificial solution-dependent estimators. The evaluation of solvers in the SimOpt library involves two main procedures. First, we run m macroreplications for each solver and problem. The solver aims to solve the problem during each macroreplication until a predefined budget is exhausted. At each $\boldsymbol{x}$, the objective function is estimated by conducting n replications using sample average approximation, which varies depending on the solver used (adaptive solvers use a random sample size $N(\boldsymbol{x})$). Second, we conduct ℓ post-replications at the intermediate incumbents of each macroreplication to estimate the objective function without optimization bias. In our experiments, we test the performance of the solvers using $m = 20$ macroreplications and $\ell = 200$ post-replications.

We use the following standard parameters: $\mu = 1000, \eta = 0.5, \gamma_1 = 0.75$, and $\gamma_2 = 1.5$. To determine $\Delta_{\max}$ for each macroreplication, we employ a process that generates random design points for the problem of interest, and the maximum distance between them is calculated and set as $\Delta_{\max}$. For each sample path, we tune Δ_0 by a pilot run as one of three possibilities, $0.05\Delta_{\max} \times (0.1, 1, 10)$ using 1% of the total budget for each. We also tune the value of the scaling parameter κ at the first iteration for each sample path by setting $\kappa = \bar{F}(\boldsymbol{X}_0, N_0)/\Delta_0^2$. Hence, κ also has three possibilities based on Δ_0. This tuning approach enables us to adjust the scaling of $\Delta_0, \Delta_{\max}$, and κ in response to the behavior of the optimization algorithm. ASTRO-DF algorithm utilizes local models with linear interpolation and implements a strategy of reusing design points from previous iterations by following the *AffPoints* algorithm presented by Wild *et al.* (2008), enabling the reuse of design points as extensively as possible. For details see (Eckman *et al.*, 2021).

As presented in Section 1, Figure 1 displays the solvability profiles for 60 problems from the SimOpt library. A solvability profile of a solver depicts the proportion of tested problems solved within a certain relative optimality gap. ASTRO-DF-C solver solved more than 80% of the problems within 30% of the budget, significantly outperforming the contenders. Next, we will examine each refinement and its corresponding effect.

5.1. Effect of diagonal Hessian

Ha *et al.* (2021) compare (i) ASTRO-DF with full Hessian, (ii) ASTRO-DF with diagonal Hessian following Definition 2.2, and (iii) ASTRO-DF that integrates both linear and fully quadratic models through a heuristic approach that utilizes linear models when far from first-order optimality and quadratic models otherwise. Experimenting with these three versions on three problems from the SimOpt library indicated that the diagonal Hessian version was capable of the fastest progress with robustness (lower variance). In the remainder of this section, we include a 20-dimensional problem to further investigate the algorithm behavior in higher dimensions.

5.2. Effect of direct search

We investigate the effect of using a direct search within ASTRO-DF after having implemented the first refinement that yields diagonal Hessian quadratic models using coordinate basis placements of interpolation points. In addition to a broad comparison of the new solver on 60 problems as illustrated in Figure 1, we conduct experiments for two problems, namely, the Stochastic Activity Network (SAN) problem, which is a convex 13-dimensional problem, and a 20-dimensional Rosenbrock function with multiplicative error,

$$F(\boldsymbol{x},\xi)=\sum_{i=1}^{19}\left[100(x_{i+1}-\xi_i x_i^2)^2+(\xi_i x_i-1)^2\right],$$

where $\xi_i \sim \mathcal{N}(1,0.1)$ for all $i \in \{1,\ldots,19\}$ (Kim and Zhang, 2010). The reason why this function has become a popular choice for evaluating optimization algorithms is attributable, in part, to the fact that its global minimum is located within a long, narrow valley that displays a parabolic shape (highly nonconvex). This characteristic makes the problem particularly challenging.

Along with Figure 2, Figure 4 now shows the solver's finite time performance in terms of the progress made per iteration with mean and 95% confidence interval after running the algorithms 20 times. In SimOpt and typical simulation experiments, a fixed total simulation budget is given; hence, the number of iterations completed varies from run to run and solver to solver. We display the first 100 iterations for the two versions of ASTRO-DF. For both problems, ASTRO-DF exhibits a slower rate of progress than ASTRO-DF-C after the first few iterations while revealing a slightly faster progress leading up to that point. This phenomenon arises from the fact that, although ASTRO-DF can attain more accurate local models by employing larger sample sizes and smaller trust regions in the first 20 or so iterations, the step size becomes excessively small thereafter. When evaluating progress relative to the expended budget, the disparity in the trajectory between the two algorithms during the initial iterations is notably less significant, as will be detailed in Figure 6.

To see per-iteration simulation expenses, see Figure 5. As shown in Figure 5(b), ASTRO-DF can only have about 70 iterations with a 30,000 budget, whereas ASTRO-DF-C reaches 100 iterations with the same budget on the SAN problem. The evident cause of less progress in ASTRO-DF than in ASTRO-DF-C is the rapid reduction of Δ_k to find new incumbents that provide satisfactory reduction. This rapid reduction of step size forces larger budget per iteration – a faster increase in the expended budget after iteration 20. Notably, the small variation in the budget spent during the initial 20 iterations is due to the fact that the deterministic lower bound of the sample size, λ_k, provides small enough standard error as in the adaptive sampling rule (6) within these initial 20 iterations.

Lastly, Figure 6 shows the mean progress with 95% confidence for each problem as a function of the expended budget. In both cases, significantly better solutions are

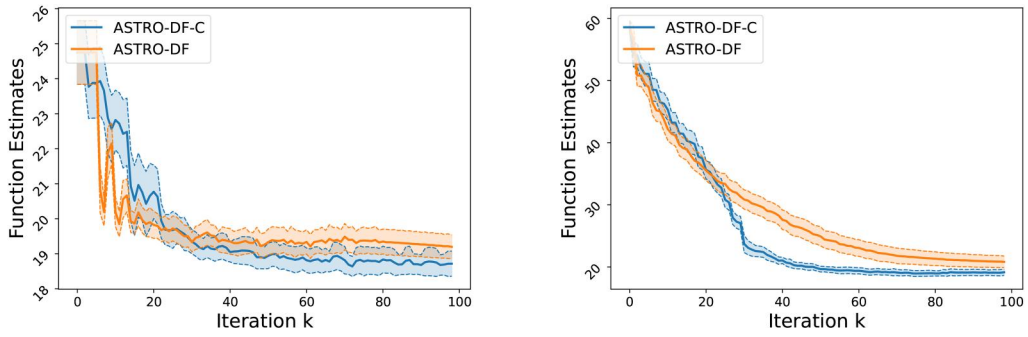

(a) 20-D Noisy Rosenbrock (b) 13-D Stochastic Activity Network

Figure 4. Better progress of ASTRO-DF-C is evident in the mean function estimates and 95% confidence intervals on both problems.

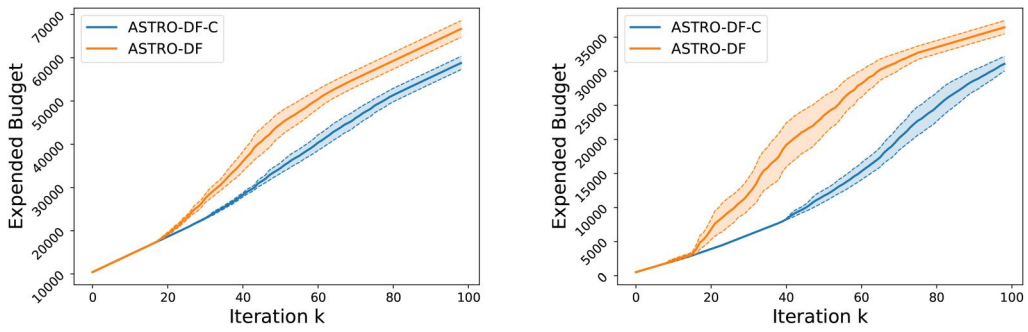

(a) 20-D Noisy Rosenbrock (b) 13-D Stochastic Activity Network

Figure 5. The total budget spent per iteration is lower due to savings from using direct search with 95% confidence intervals from 20 macroreplications.

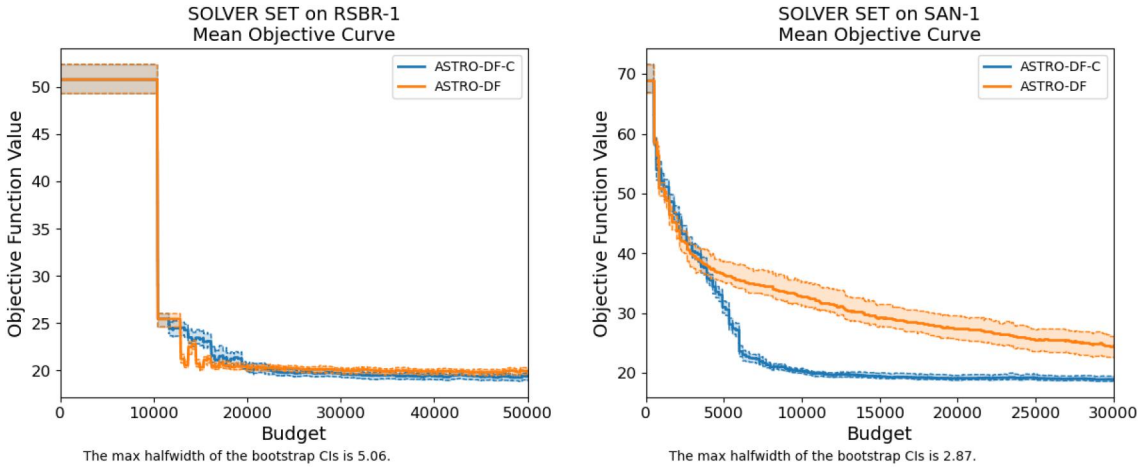

(a) 20-D Noisy Rosenbrock (b) 13-D Stochastic Activity Network

Figure 6. Mean value and 95% confidence interval of objective function trajectory per spent budget for both problems exhibits significant finite-time improvement with direct search.

reached by ASTRO-DF-C. Especially for the SAN problem, we observe that direct search remarkably accelerates the convergence with less uncertainty (narrower confidence intervals). In the Rosenbrock function, the improvement, while significant, is not as pronounced. We attribute consistently better solutions with direct search in this non-convex noisy problem to a better exploration of the feasible region (more chances of success means visiting more points) with no additional cost, besides the saving of budget and a slower rate of Δ_k decay.

6. Conclusion and future work

This article proposes efficient procedures and theoretical advancements for solving SDFO to first-order optimality in finite time. Trust regions are a class of algorithms with growing popularity for their volatile mechanics and stability in this context. However, expensive function approximations with locally certified models have caused a shortage of implementable trust-region algorithms for SDFO that serve higher dimensions robustly and efficiently. Our approach of integrating adaptive sampling strategies with local approximations via coordinate-basis designs and leveraging intermediate direct search steps addresses this limitation. The resulting algorithm is implemented and experimented on a testbed of SDFO problems and maintains almost sure convergence guarantees. Crucially, we also prove canonical complexity rates in almost sure sense and in expectation without requiring practically hard assumptions. The gain in finite-time performance is apparent with justifiable faster progress in the early iterations of the search.

A future research direction is on the exploration of the work complexity or total oracle runs and the dimension-dependent constants. Additionally, reusing history and allowing the design set geometry to vary may be beneficial. Although ASTRO-DF-C effectively utilizes information by incorporating direct search, the data (visited points) from previous iterations is not carried over to subsequent iterations. Therefore, there might be potential to enhance efficiency by systematically reusing the design points and replications while retaining the benefits of ASTRO-DF-C. Other research areas include handling higher dimensions with techniques such as random subspaces, in which one generates a sequence of incumbents to random embedding constructions in lower dimensions without losing much information (Dzahini and Wild, 2022; Cartis and Roberts, 2023).

Acknowledgments

The authors acknowledge the helpful discussions with Dr. Raghu Pasupathy that led to the improvement of the manuscript.

Funding

Generous support from the National Science Foundation Grant CMMI-2226347 is gratefully acknowledged.

Notes on contributors

Dr. Yunsoo Ha is currently a Postdoctoral Researcher at the National Renewable Energy Laboratory within the Artificial Intelligence, Learning, and Intelligent Systems (ALIS) group. He earned his PhD degree from Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. His research revolves around Machine Learning, Quantum Computing, and Stochastic Optimization. More precisely, his focus lies in crafting stochastic optimization algorithms tailored to address real-world challenges.

Dr. Sara Shashaani is an Assistant Professor and Bowman Faculty Scholar in Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. Her research interests lie in the intersection of Monte Carlo simulation, applied probability, and stochastic optimization. She is an active member of the INFORMS Simulation Society and INFORMS Computing Society. She is also the co-creator of SimOpt: the open-source testbed and benchmarking platform for simulation optimization algorithms.

ORCID

Yunsoo Ha http://orcid.org/0000-0002-9421-4768
Sara Shashaani http://orcid.org/0000-0001-8515-5877

Data availability statement

The data and computer code that support the findings of this study are openly available in GitHub at https://github.com/sshashaa/astro-df.git.

References

Audet, C., Digabel, S.L., Montplaisir, V.R. and Tribes, C. (2021) NOMAD version 4: Nonlinear optimization with the MADS algorithm. *arXiv preprint arXiv:2104.11627*.

Berahas, A.S., Cao, L., Choromanski, K. and Scheinberg, K. (2021) A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Foundations of Computational Mathematics*, **22**(2), 507–560.

Berahas, A.S., Cao, L. and Scheinberg, K. (2021) Global convergence rate analysis of a generic line search algorithm with noise. *SIAM Journal on Optimization*, **31**(2), 1489–1518.

Blanchet, J., Cartis, C., Menickelly, M. and Scheinberg, K. (2019) Convergence rate analysis of a stochastic trust-region method via supermartingales. *INFORMS Journal on Optimization*, **1**(2), 92–119.

Cakmak, S., Wu, D. and Zhou, E. (2021) Solving Bayesian risk optimization via nested stochastic gradient estimation. *IISE Transactions*, **53**(10), 1081–1093.

Cao, L., Berahas, A.S. and Scheinberg, K. (2022) First- and second-order high probability complexity bounds for trust-region methods with noisy oracles. *arXiv:2205.03667*.

Cartis, C. and Roberts, L. (2023) Scalable subspace methods for derivative-free nonlinear least-squares optimization. *Mathematical Programming*, **199**, 461–524.

Chang, K.-H., Hong, L.J. and Wan, H. (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., Menickelly, M. and Scheinberg, K. (2018) Stochastic optimization using a trust-region method and random models. *Mathematical Programming*, **169**(2), 447–487.

Choromanski, K., Iscen, A., Sindhwani, V., Tan, J. and Coumans, E. (2018) Optimizing simulations with noise-tolerant structured exploration, in *2018 IEEE International Conference on Robotics and Automation*, IEEE Press, Piscataway, NJ, pp. 2970–2977.

Conn, A.R., Scheinberg, K. and Vicente, L.N. (2009) *Introduction to Derivative-Free Optimization*. 1st edition, SIAM, Philadelphia, PA.

Coope, I.D. and Tappenden, R. (2021) Gradient and diagonal Hessian approximations using quadratic interpolation models and aligned regular bases. *Numerical Algorithms*, **88**, 767–791.

Dzahini, K.J. and Wild, S.M. (2022) Stochastic trust-region algorithm in random subspaces with convergence and expected complexity analyses. *arXiv.2207.06452*.

Eckman, D.J., Henderson, S.G. and Shashaani, S. (2023) Diagnostic tools for evaluating and comparing simulation-optimization algorithms. *INFORMS Journal on Computing*, **35**(2), 350–367.

Eckman, D.J., Henderson, S.G., Shashaani, S. and Pasupathy, R. (2021) SimOpt. https://github.com/simopt-admin/simopt (accessed 1 May 2020).

Fazel, M., Ge, R., Kakade, S. and Mesbahi, M. (2018) Global convergence of policy gradient methods for the linear quadratic regulator, in *Proceedings of the 35th International Conference on Machine Learning*, PMLR (Proceedings of Machine Learning Research), Stockholm, Sweden, pp. 1467–1476.

Flaxman, A.D., Kalai, A.T. and McMahan, H.B. (2004) Online convex optimization in the bandit setting: Gradient descent without a gradient. arXiv:cs/0408007.

Ghanbari, H. and Scheinberg, K. (2017) Black-box optimization in machine learning with trust region based derivative free algorithm. *arXiv.1703.06925*.

Gratton, S., Royer, C.W., Vicente, L.N. and Zhang, Z. (2018) Complexity and global rates of trust-region methods based on probabilistic models. *IMA Journal of Numerical Analysis*, **38**(3), 1579–1597.

Ha, Y., Shashaani, S., and Pasupathy, R. (2023). On common random numbers and the complexity of adaptive sampling trust-region methods. *optimization-online.org/?p=*23853.

Ha, Y., Shashaani, S. and Tran-Dinh, Q. (2021) Improved complexity of trust-region optimization for zeroth-order stochastic oracles with adaptive sampling, in *Proceedings of the 2021 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 1–12.

Hu, J., Song, M. and Fu, M.C. (2023) Quantile optimization via multiple timescale local search for black-box functions. *arXiv preprint arXiv:2308.07607*.

Jin, B., Scheinberg, K. and Xie, M. (2021) High probability complexity bounds for line search based on stochastic oracles. in *Advances in Neural Information Processing Systems*, **34**, 9193–9203.

Kim, S. and Zhang, D. (2010) Convergence properties of direct search methods for stochastic optimization, in *Proceedings of the 2010 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 1003–1011.

Kingma, D.P. and Ba, J. (2017) Adam: A method for stochastic optimization. *arXiv:1412.6980*.

Ko, B. and Tang, Q. (2008) Sums of dependent nonnegative random variables with subexponential tails. *Journal of Applied Probability*, **45**(1), 85–94.

L'Ecuyer, P., Simard, R., Chen, E.J. and Kelton, W.D. (2002) An object-oriented random number package with many long streams and substreams. *Operations Research*, **50**(6), 1073–1075.

Maggiar, A., Wachter, A., Dolinskaya, I.S. and Staum, J. (2018) A derivative-free trust-region algorithm for the optimization of functions smoothed via Gaussian convolution using adaptive multiple importance sampling. *SIAM Journal on Optimization*, **28**(2), 1478–1507.

Menickelly, M., Ha, Y. and Otten, M. (2023) Latency considerations for stochastic optimizers in variational quantum algorithms. *Quantum*, **7**, 949.

Paquette, C. and Scheinberg, K. (2020) A stochastic line search method with expected complexity analysis. *SIAM Journal on Optimization*, **30**(1), 349–376.

Ragonneau, T.M. and Zhang, Z. (2023) An optimal interpolation set for model-based derivative-free optimization methods. *arXiv:2302.09992*.

Ruan, Y., Xiong, Y., Reddi, S., Kumar, S. and Hsieh, C.-J. (2020) Learning to learn by zeroth-order oracle. *arXiv:1910.09464*.

Salimans, T., Ho, J., Chen, X., Sidor, S. and Sutskever, I. (2017) Evolution strategies as a scalable alternative to reinforcement learning. *arXiv:1703.03864*.

Shashaani, S., Hunter, S.R. and Pasupathy, R. (2016) ASTRO-DF: Adaptive sampling trust-region optimization algorithms, heuristics, and numerical experience, in *Proceedings of the 2016 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 554–565.

Shashaani, S., Hashemi, F.S., and Pasupathy, R. (2018) ASTRO-DF: A class of adaptive sampling trust-region algorithms for derivative-free stochastic optimization. *SIAM Journal on Optimization*, **28**(4), 3145–3176.

Wild, S.M., Regis, R.G. and Shoemaker, C.A. (2008) ORBIT: Optimization by radial basis function interpolation in trust-regions. *SIAM Journal on Scientific Computing*, **30**(6), 3197–3219.

Xu, J. and Zheng, Z. (2023) Gradient-based simulation optimization algorithms via multi-resolution system approximations. *INFORMS Journal on Computing*, **35**(3), 633–651.

Xu, W.L. and Nelson, B.L. (2013) Empirical stochastic branch-and-bound for optimization via simulation. *IIE Transactions*, **45**(7), 685–698.

Zhang, D., Ji, L., Zhao, S. and Wang, L. (2023) Variable-sample method for the computation of stochastic Nash equilibrium. *IISE Transactions*, **55**(12), 1217–1229.