Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

IMPROVED COMPLEXITY OF TRUST-REGION OPTIMIZATION FOR ZEROTH-ORDER STOCHASTIC ORACLES WITH ADAPTIVE SAMPLING

Yunsoo Ha
Sara Shashaani

Edward P. Fitts Department of
Industrial and Systems Engineering
North Carolina State University
915 Partners Way
Raleigh, NC 27695, USA

Quoc Tran-Dinh

Department of Statistics
and Operations Research
University of North Carolina at Chapel Hill
333 Hanes Hall
Chapel Hill, NC 27599, USA

ABSTRACT

We present an enhanced stochastic trust-region optimization with adaptive sampling (ASTRO-DF) in which optimizing an iteratively constructed local model on estimates of objective values with stochastic sample size guides the search. The noticeable feature is that the underdetermined quadratic model with a diagonal Hessian requires fewer function evaluations, which is particularly useful at high dimensions. This paper describes the enhanced algorithm in detail. It gives several theoretical results, including iteration complexity, and renders almost sure convergence guarantees. We report in our numerical experience the finite-time superiority of the enhanced ASTRO-DF over state-of-the-art using the SimOpt library.

1 INTRODUCTION

This work considers unconstrained stochastic optimization (SO) problems of the form

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

where $f:\mathbb{R}^d\to\mathbb{R}$ is smooth and bounded from below, and $F:\mathbb{R}^d\times\Omega\to\mathbb{R}$ is defined on a probability space $(\Omega,\mathcal{F},P)$. We denote the optimal value by $f^* := \inf_{\boldsymbol{x}\in\mathbb{R}^d} f(\boldsymbol{x}) > -\infty$. Since $f(\boldsymbol{x})$ is only observable by a Monte Carlo simulation, we can generate the random variable $F(\boldsymbol{x})$. Hence, the estimator of $f(\boldsymbol{x})$ can be obtained by $f(\boldsymbol{x},n) = n^{-1}\sum_{i=1}^n F(\boldsymbol{x},\xi_i)$, and let $\hat{\sigma}_F^2(\boldsymbol{x},n) = (n-1)^{-1}\sum_{j=1}^n(F(\boldsymbol{x},\xi_i)-f(\boldsymbol{x},n))^2$ be its estimated variance. Furthermore, the consider zeroth-order stochastic oracles, where the derivative information is not directly available from the Monte Carlo Simulation.

Trust region optimization (TRO) is a popular solver for nonlinear and nonconvex problems, especially in noisy settings. Recently, developing stochastic trust-region methods has gained a lot of attention. STRONG (Chang et al. 2013) obtains first-order convergence almost surely under certain conditions for the problem and the estimators for first-order oracles. STORM (Chen et al. 2018) is another stochastic TRO algorithm that uses random models of specified accuracy.

SO on zeroth-order oracles, also known as derivative-free (DF) optimization, is particularly difficult for its rapidly growing cost with dimension. A key component of DF algorithms is their implicit or explicit estimation of the underlying function's gradient and Hessian at each iteration. We can approximate the derivative information through several methods: finite-difference approximations, response surface methodology via interpolation models, and Gaussian smoothing. We focus on interpolation and how it can be made more efficient for higher-dimensional problems. TRO-DF via interpolation requires function evaluations at several points in a neighborhood of the incumbent solution. In particular, to construct a

978-1-6654-3311-2/21/$31.00 ©2021 IEEE

quadratic model with interpolation, the function needs to be estimated at $(d+1)(d+2)/2$ number of points. However, Monte Carlo oracle calls often have a high computational burden. This requirement quickly becomes an impediment for efficiency, not only because, especially in stochastic settings, estimation error at each point adds to the source of error, but also because picking the best set of points in a higher dimension quickly becomes cumbersome. ASTRO-DF (Shashaani et al. 2018) devises an algorithm that addresses efficiency by selecting the sample size, i.e., the number of times the oracle is invoked at each point adaptively. The main idea of adaptive sampling is to limit the effort commensurate with inferred optimality gap: far solutions need fewer samples while closer solutions need more. Despite promising theoretical and empirical results, ASTRO-DF still suffers from the curse of dimensionality.

Our paper combines two main ideals to design a new derivative-free optimization solver for (1). The first idea is the ASDTRO-DF scheme with an improved sample size lower bound. The second one is a derivative estimation via interpolation using the coordinate basis by Coope and Tappenden (2020) that requires only $\mathcal{O}(d)$ number of points. Utilizing these ideas, we propose a new variant of ASTRO-DF that has convergence rate guarantees which its original version lacks. The improvements are shown theoretically in Theorem 3 and empirically in Section 4. Our theoretical results are the first to prove almost sure iteration complexity; the existing optimization algorithms for stochastic simulations only guarantee iteration complexity in expectation (Blanchet et al. 2019).

1.1 ASTRO-DF Algorithm

ASTRO-DF follows the same logic as its deterministic counterpart (DTRO-DF) algorithms (Shashaani et al. 2018). At each iteration k, the objective function is improved by optimizing a second-order model constructed only via function value estimates at the incumbent solution $\boldsymbol{X}_k$ and $d(d+3)/2$ design points with a poised spread around $\boldsymbol{X}_k$. At each point, the number of Monte Carlo oracle calls is adaptive by balancing the standard error with the optimality error measured by the size of the trust-region Δ_k to the fourth power. Furthermore, special care ensures that the quadratic model's gradient remains in lock-step with Δ_k. After the model construction, minimizing the local model (inexactly) within the trust region obtains the next candidate incumbent $\tilde{\boldsymbol{X}}_{k+1}$ for iteration $k+1$. $\tilde{\boldsymbol{X}}_{k+1}$ is accepted, and the trust-region expanded if the success ratio, which reflects how well the model has predicted the underlying function, is large enough. Otherwise, a new candidate incumbent solution is attempted in the next iteration from the same place but with points selected from a contracted trust region. ASTRO-DF globally converges to a first-order critical point almost surely. Global convergence here means that irrespective of where the starting solution is, the algorithm is guaranteed to converge almost surely.

1.2 Notations and Definitions

We use bold font for vectors; hence $\boldsymbol{x} = (x_1, x_2, \cdots, x_d) \in \mathbb{R}^d$ denotes a d-dimensional vector of real numbers. Let $\boldsymbol{e}_i \in \mathbb{R}^d$ for $i = 1, \ldots, d$ denote the standard unit basis vectors in $\mathbb{R}^d$, and $\mathcal{B}(\boldsymbol{y}^0; \delta) = \{\boldsymbol{y} \in \mathbb{R}^d : \|\boldsymbol{y} - \boldsymbol{y}^0\|_2 \leq \delta\}$ be the closed ball of radius $\delta > 0$ centered at $\boldsymbol{y}^0$. $\mathcal{Y} = \{\boldsymbol{y}^0, \boldsymbol{y}^1, \ldots, \boldsymbol{y}^p\} \subset \mathcal{B}(\boldsymbol{y}^0; \delta)$ denotes the sample set which is used to fit a local model that can help approximate the gradient and Hessian. Let $t_\epsilon := \min\{k \in \mathbb{N} : \|\nabla f(\boldsymbol{x}_k)\| \leq \epsilon\}$ be the first-order ϵ-stationary stopping time.

Next, we introduce several definitions which will be invoked during the rest of the paper.

Definition 1 (polynomial interpolation models). Given $\boldsymbol{y} \in \mathbb{R}^d$ and $\delta > 0$, let $f : \mathbb{R}^d \to \mathbb{R}$ be a real-valued function and $\Phi(\boldsymbol{z}) = (\phi_0(\boldsymbol{z}), \phi_1(\boldsymbol{z}), \ldots, \phi_q(\boldsymbol{z}))$ be a polynomial basis on $\mathbb{R}^d$. With $p = q$ and $\mathcal{Y}$, suppose that we can find $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \ldots, \alpha_p)^\intercal$ such that $\mathcal{M}(\Phi, \mathcal{Y})\boldsymbol{\alpha} = f(\mathcal{Y})$, where

$$\mathcal{M}(\Phi, \mathcal{Y}) = \begin{bmatrix} \phi_0(\boldsymbol{y}^0) & \phi_1(\boldsymbol{y}^0) & \cdots & \phi_q(\boldsymbol{y}^0) \\ \phi_0(\boldsymbol{y}^1) & \phi_1(\boldsymbol{y}^1) & \cdots & \phi_q(\boldsymbol{y}^1) \\ \vdots & \vdots & \vdots & \vdots \\ \phi_0(\boldsymbol{y}^p) & \phi_1(\boldsymbol{y}^p) & \cdots & \phi_q(\boldsymbol{y}^p) \end{bmatrix}, \quad \boldsymbol{\alpha} = \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_p \end{bmatrix}, \quad f(\mathcal{Y}) = \begin{bmatrix} f(\boldsymbol{y}^0) \\ f(\boldsymbol{y}^1) \\ \vdots \\ f(\boldsymbol{y}^p) \end{bmatrix}.$$

Then, the function $m(\boldsymbol{z}) : \mathcal{B}(\boldsymbol{y}^0;\delta) \to \mathbb{R}$, as $m(\boldsymbol{z}) = \sum_{j=0}^{p} \alpha_j \phi_j(\boldsymbol{z})$ is a polynomial interpolation model of f on $\mathcal{B}(\boldsymbol{y}^0;\delta)$. Let $g(\boldsymbol{y}^0) = \begin{bmatrix}\alpha_1 & \alpha_2 & \cdots & \alpha_d\end{bmatrix}^\intercal$ be the subvector of $\boldsymbol{\alpha}$ and $\mathcal{H}(\boldsymbol{y}^0)$ be a symmetric matrix of order d. The elements of $\mathcal{H}(\boldsymbol{y}^0)$ are uniquely defined by $\begin{bmatrix}\alpha_{d+1} & \alpha_{d+2} & \cdots & \alpha_p\end{bmatrix}^\intercal$. Then, we can define a quadratic interpolating polynomial model $m : \mathcal{B}(\boldsymbol{y}^0;\delta) \subset \mathbb{R}^d \to \mathbb{R}$, as

$$m(\boldsymbol{y}) = m(\boldsymbol{y}^0) + (\boldsymbol{y}-\boldsymbol{y}^0)^\intercal g(\boldsymbol{y}^0) + \frac{1}{2}(\boldsymbol{y}-\boldsymbol{y}^0)^\intercal \mathcal{H}(\boldsymbol{y}^0)(\boldsymbol{y}-\boldsymbol{y}^0), \tag{2}$$

where $\boldsymbol{y}^0 \in \mathbb{R}^d$ and $\delta > 0$.

Definition 2 (diagonal quadratic interpolation models with coordinate bases) A special case of (2) is when the Hessian has only diagonal values, which we denote by $\mathcal{D}(\boldsymbol{y}^0)$ and define as

$$\mathcal{D}(\boldsymbol{y}^0) = \begin{bmatrix} h_1 & & \\ & \ddots & \\ & & h_d \end{bmatrix} \in \mathbb{R}^{d\times d}.$$

In the quadratic interpolation model with diagonal Hessian $p = 2d$ and the model defined in (2) contains $2d+1$ unknowns. Hence, $2d+1$ interpolation points and function evaluations are needed to uniquely determine the $g(\boldsymbol{y}^0)$ and $\mathcal{D}(\boldsymbol{y}^0)$. We set

$$\mathcal{Y}_{cb} = \{\boldsymbol{y}^0, \boldsymbol{y}^0 + \boldsymbol{e}_1\delta, \ldots, \boldsymbol{y}^0 + \boldsymbol{e}_d\delta, \boldsymbol{y}^0 - \boldsymbol{e}_1\delta, \ldots, \boldsymbol{y}^0 - \boldsymbol{e}_d\delta\}$$

contained in the ball $\mathcal{B}(\boldsymbol{y}^0;\delta)$. Since the coordinate basis is used to generate the interpolation points, $\boldsymbol{\alpha}$ is guaranteed to exist. Hence, $h_i \le \infty$ for all $i = 1, 2, \ldots, d$. In this case,

$$\bar{\Phi}(\boldsymbol{z}) := (1, z_1, z_2, \ldots, z_d, z_1^2, z_2^2, \ldots, z_d^2),$$

and $m(\boldsymbol{z})$ is said to be a diagonal quadratic interpolation model.

Definition 3 (fully linear models) Given $\boldsymbol{y} \in \mathbb{R}^d$ and $\delta > 0$, a model function $m(\boldsymbol{y}) : \mathcal{B}(\boldsymbol{y}^0;\delta) \to \mathbb{R}$ is the fully linear model of f if $\nabla m(\boldsymbol{y})$ is Lipschitz continuous with constant ν_1^m, and there exist two positive constants κ_{eg} and κ_{ef} such that

$$\begin{aligned}\|\nabla f(\boldsymbol{y}) - \nabla m(\boldsymbol{y})\| &\le \kappa_{eg}\delta,\\ \|f(\boldsymbol{y}) - m(\boldsymbol{y})\| &\le \kappa_{ef}\delta^2.\end{aligned}$$

Definition 4 (Cauchy reduction) Given $m(\cdot)$ on $\mathcal{B}(\boldsymbol{x};\delta)$, $\boldsymbol{s}^c$ is the Cauchy step if

$$m(\boldsymbol{x}) - m(\boldsymbol{x}+\boldsymbol{s}^c) \ge \frac{1}{2}\|\nabla m(\boldsymbol{x})\| \min\left\{\frac{\|\nabla m(\boldsymbol{x})\|}{\|\nabla^2 m(\boldsymbol{x})\|}, \delta\right\}.$$

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 $\mathcal{B}(\boldsymbol{x};\delta)$.

The standing assumptions for the remainder of the paper are:

Assumption 1 We assume that the function f is twice continuously differentiable in an open domain Ω containing $\mathcal{B}(\boldsymbol{y}^0;\delta)$, ∇f is Lipschitz continuous in Ω with Lipschitz constant $\kappa_{Lg} > 0$.

Assumption 2 There exists a constant $\kappa_{fcd} \in (0, 1]$ such that for all iterations k we have

$$m_k(\boldsymbol{x}_k) - m_k(\boldsymbol{x}_k + s_k) \ge \kappa_{fcd}[m_k(\boldsymbol{x}_k) - m_k(\boldsymbol{x}_k + s_k^c)],$$

where s_k^c is the Cauchy step.

2 EFFECTS OF UNDER-DETERMINED LOCAL MODEL ON DETERMINISTIC TRO-DF

Before analyzing the complexity of the improved ASTRO-DF, we first provide the iteration complexity and the oracle complexity for DTRO-DF with diagonal Hessian. On iteration k, DTRO-DF with diagonal Hessian follows the same steps of Algorithm 1 in (Shashaani et al. 2016), except that the interpolation set for the local model construction is selected based on the coordinate basis, that is, in a deterministic manner as described in Definition 2. Consequently, the quadratic model is under-determined.

2.1 Consistency

DTRO-DF converges to a first-order critical point globally as long as the local model is fully linear (see Definition 3). We show that this consistency result is not lost by proving that the local model with diagonal Hessian is fully linear model. The general underdetermined interpolating model can easily be shown to follow the full-linearity guarantee in Theorem 5.4 (Conn et al. 2009b). More specifically we can trivially show the following result and choose not to include the proof due to space limits.

Theorem 1 Let Assumption 1 hold, and the interpolation model $m(\boldsymbol{z})$ be constructed via points selected following a coordinated basis, i.e., $\mathcal{Y}_{cb}$. Given that the model gradient $\nabla m(\boldsymbol{y}) = (\boldsymbol{y} - \boldsymbol{y}^0)^\intercal \mathcal{D}(\boldsymbol{y}^0) + g(\boldsymbol{y}^0)$ for all points $\boldsymbol{y}$ in $\mathcal{B}(\boldsymbol{y}^0; \delta)$, let $0 < \kappa_{\mathcal{D}}$ be the uniform upper bound on the model Hessian norm $\|\mathcal{D}(\boldsymbol{y}^0)\|$. Then, we can uniformly bound the model gradient error by

$$\|\nabla m(\boldsymbol{y}) - \nabla f(\boldsymbol{y})\| \leq \kappa_{eg}\delta,$$

where $\kappa_{eg} = \frac{5\sqrt{2d}}{2}(\kappa_{Lg} + \kappa_{\mathcal{D}})$.

From Theorem 1, the next result which is the same as Theorem 10.12 (Conn et al. 2009b) holds true.

Theorem 2 Suppose $\{\boldsymbol{x}_k\}$ is a sequence generated by the DTRO-DF Algorithm. Let Assumption 1 and 2 hold and $\kappa_{\mathcal{D}} > 0$ be the uniform upper bound on the model Hessian norm $\|\mathcal{D}_k(\boldsymbol{x}_k)\|$ for all k. Moreover, suppose that the model $m_k(\cdot)$ is fully linear on $\mathcal{B}(\boldsymbol{x}_k; \check{\delta}_k)$ with model gradient error constant κ_{eg}. Then,

$$\liminf_{k \to +\infty} \|\nabla f(\boldsymbol{x}_k)\| = 0.$$

We notice that $\lim_{k\to+\infty} \|\nabla f(\boldsymbol{x}_k)\| = 0$ is also satisfied from Theorem 2 by the continuity of f. Due to space limit we refer the interested reader to (Shashaani et al. 2018) for more details.

2.2 Iteration and Oracle Complexity of the DTRO-DF Algorithm

We know that the DTRO-DF Algorithm converges to the first-order critical point. The first-order ϵ-stationary stopping time t_ϵ can be obtained by the existing result (Curtis and Scheinberg 2020), listed below.

Theorem 3 Let $\{\tau_k\}$ be a sequence of nonnegative constants, $l_\epsilon : [0, \infty) \to (0, \infty)$ be a non-decreasing function and $\Theta \in (0, \infty)$ be a scalar such that, for all $k < t_\epsilon$, if iteration k is successful, i.e., $\rho_k = \frac{f(\boldsymbol{x}_k) - f(\tilde{\boldsymbol{x}}_{k+1})}{m_k(\boldsymbol{x}_k) - m_k(\tilde{\boldsymbol{x}}_{k+1})} \geq \eta_1$ for a given $\eta_1 > 0$, $\tau_k - \tau_{k+1} \geq \Theta l_\epsilon(\delta_k)$, and δ_ϵ be a non-negative constant such that for all $k \leq t_\epsilon$, $\delta_k \leq \delta_\epsilon$ implies iteration k is successful. Then $t_\epsilon \leq \mathcal{O}\left(\frac{\tau_0}{l_\epsilon(\delta_\epsilon)}\right)$.

To obtain the oracle complexity, we need the iteration complexity and the arithmetic complexity for each iteration. The following new result ensures that the contraction loop (see Algorithm 1 for its stochastic version) terminates with in a finite number of steps that is bounded by a bound that is independent of the iteration.

Theorem 4 In the DTRO-DF Algorithm, for all $k < t_\epsilon$ given some $\epsilon > 0$, the contraction loop terminates in uniformly bounded finitely many steps where the uniform upper bound is

$$\kappa_{clu} = \left\lceil \log_w \frac{\mu\epsilon}{(\mu\kappa_{eg} + 1)\delta_{max}} + 1 \right\rceil.$$

Proof. For an arbitrary iteration $k < t_\epsilon$ for which we have $\|\nabla f(\boldsymbol{x}_k)\| \geq \epsilon$, let j_k be the last run of the contraction loop, i.e., $\|\nabla m_k^{(j)}(\boldsymbol{x}_k)\| < \mu^{-1}\delta_k^{(j)}$ for all $j < j_k$. Consider $j_k - 1$, one to the last run of the contraction loop at iteration k. Full-linearity of the model $m_k(\cdot)$ satisfies

$$\|\nabla f(\boldsymbol{x}_k) - \nabla m_k^{(j_k-1)}(\boldsymbol{x}_k)\| \leq \kappa_{eg}\delta_k w^{j_k-2}.$$

Then, by the triangle inequality

$$\begin{aligned}\|\nabla m_k^{(j_k-1)}(\boldsymbol{x}_k)\| &\geq \|\nabla f(\boldsymbol{x}_k)\| - \|\nabla f(\boldsymbol{x}_k) - \nabla m_k^{(j_k-1)}(\boldsymbol{x}_k)\| \\ &\geq \epsilon - \kappa_{eg}\delta_k w^{j_k-2}.\end{aligned}$$

Hence, we get $\epsilon - \kappa_{eg}\delta_k w^{j_k-2} < \mu^{-1}\delta_k w^{j_k-2}$ from which we obtain

$$\begin{aligned} j_k &< \frac{\log(\mu\epsilon) - \log((1+\mu\kappa_{eg})\delta_{\max})}{\log w} + 2 \\ &= \frac{\log(\mu\epsilon) - \log(1+\mu(\frac{5\sqrt{2d}}{2}(\kappa_{Lg}+\kappa_{\mathcal{D}})))\delta_{\max}}{\log w} + 2, \end{aligned} \tag{3}$$

where the second equality comes from Theorem 1. Consequently, $j_k \leq \kappa_{clu}$ for all $k < t_\epsilon$. □

Notice from (3) that we can say $\kappa_{clu} = \mathcal{O}\left(\log \frac{d}{\epsilon}\right)$. We will be using the same uniform bound in the stochastic proofs in Section 3, by later proving that the additional stochastic error is notwithstanding the results. To obtain the arithmetic complexity at iteration k, given the uniform upper bound on the number of function evaluations at iterations k, we need to know how many iterations are required to achieve an ϵ-stationary solution.

Theorem 5 Let $\{\boldsymbol{x}_k\}$ be a sequence of iterates generated by the DTRO-DF Algorithm, and let Assumptions 1 and 2 hold. Then the first-order ϵ-stationary stopping time satisfies $t_\epsilon = \mathcal{O}(\epsilon^{-2})$.

Proof. In order to prove Theorem 5, we need to show that DTRO-DF satisfies two conditions mentioned in Theorem 3. We define $\tau_k := v(f(\boldsymbol{x}_k) - f^*) + (1-v)\delta_k^2$ for some $v \in (0,1)$. Then, we get

$$\tau_k - \tau_{k+1} = v(f(\boldsymbol{x}_k) - f(\boldsymbol{x}_{k+1})) + (1-v)(\delta_k^2 - \delta_{k+1}^2).$$

By Assumption 2, if iteration k is very successful, we obtain

$$\begin{aligned} f(\boldsymbol{x}_k) - f(\boldsymbol{x}_{k+1}) &\geq \eta_2(m(\boldsymbol{x}_k) - m(\boldsymbol{x}_{k+1})) \\ &\geq \eta_2\frac{\kappa_{fcd}}{2}\|\nabla m_k(\boldsymbol{x}_k)\| \min\left\{\frac{\|\nabla m_k(\boldsymbol{x}_k)\|}{\kappa_{\mathcal{D}}}, \check{\delta}_k\right\} \\ &\geq \eta_2\frac{\kappa_{fcd}}{2\mu}\check{\delta}_k^2 \min\left\{\frac{1}{\kappa_{\mathcal{D}}\mu}, 1\right\} \\ &\geq \eta_2\frac{\kappa_{fcd}}{2\mu}\min\left\{\frac{1}{\kappa_{\mathcal{D}}\mu}, 1\right\} w^{2(\kappa_{clu}-1)}\delta_k^2, \end{aligned}$$

where the third inequality comes from the fact that $\check{\delta}_k \leq \mu\|\nabla m_k(\boldsymbol{x}_k)\|$. Since $\delta_{k+1} = \min\{\gamma_2\check{\delta}_k, \delta_{\max}\} \leq \gamma_2\delta_k$, we obtain

$$\begin{aligned} \tau_k - \tau_{k+1} &\geq \left(v\eta_2\frac{\kappa_{fcd}}{2\mu}\min\left\{\frac{1}{\mu\kappa_{\mathcal{D}}}, 1\right\} w^{2(\kappa_{clu}-1)}\right)\delta_k^2 + (1-v)(\delta_k^2 - \delta_{k+1}^2) \\ &\geq \left(v\eta_2\frac{\kappa_{fcd}}{2\mu}\min\left\{\frac{1}{\mu\kappa_{\mathcal{D}}}, 1\right\} w^{2(\kappa_{clu}-1)} + (1-v)(1-\gamma_2^2)\right)\delta_k^2. \end{aligned}$$

If iteration k is successful, $\tau_k - \tau_{k+1} = (1-v)(\delta_k^2 - \delta_{k+1}^2) \geq 0$, since $\delta_{k+1} = \check{\delta}_k \leq \delta_k$. Otherwise, if iteration k is unsuccessful,

$$\tau_k - \tau_{k+1} = (1-v)(\delta_k^2 - \delta_{k+1}^2) \geq (1-v)(1-\gamma_1^2)\delta_k^2 \geq 0,$$

where the last inequality comes from the fact that $\delta_{k+1} = \gamma_1 \check{\delta}_k \leq \gamma_1 \delta_k$. Then, we can see that $\tau_k - \tau_{k+1}$ is always positive by choosing v sufficiently close to 1. As a result, for $k \leq t_\epsilon$, we obtain $\tau_k - \tau_{k+1} \geq \theta \delta_k^2$, where $\theta \in (0, \infty)$. Now, we derive that there exists δ_ϵ such that $\delta_k \geq \delta_\epsilon$ for all $k \leq t_\epsilon$. For all $k \leq t_\epsilon$, we have $\|\nabla f(\boldsymbol{x}_k)\| > \epsilon$ and thus, we get $\|\nabla m_k(\boldsymbol{x}_k)\| > \|\nabla f(\boldsymbol{x}_k)\| - \|\nabla m_k(\boldsymbol{x}_k) - \nabla f(\boldsymbol{x}_k)\| > \epsilon - \kappa_{eq}\check{\delta}_k$.

Then, since $\check{\delta}_k \geq \beta \|\nabla m_k(\boldsymbol{x}_k)\|$ and $\delta_{k+1} \geq \gamma_1 \check{\delta}_k$, we obtain $\delta_{k+1} \geq \gamma_1 \check{\delta}_k \geq \frac{\gamma_1 \beta \epsilon}{1+\beta\kappa_{eg}} = \delta_\epsilon$. It means that δ_k is always greater than or equal to δ_ϵ for all $k \leq t_\epsilon$ and thus, when $\delta_k < \delta_\epsilon$, k-th iteration is always successful. Thus, the two conditions mentioned in Theorem 3 are satisfied and we get

$$t_\epsilon \leq \frac{v(f(\boldsymbol{x}_0) - f^*) + (1-v)\delta_0^2}{\theta \delta_\epsilon^2} = \mathcal{O}(\epsilon^{-2}).$$

□

Corollary 6 If we let w_ϵ be the total work, that is, the total number of function evaluations until the ϵ-accuracy stopping time t_ϵ, then $w_\epsilon = \mathcal{O}(\frac{d}{\epsilon^2}\log(\frac{d}{\epsilon}))$.

Proof. Since $2d$ new points are evaluated at each run of the model construction inner loop, the total number of function evaluations for each iteration is at most $2d\kappa_{clu} + 1$. Hence, the oracle complexity becomes $w_\epsilon \leq t_\epsilon(2d\kappa_{clu} + 1) = \mathcal{O}(\frac{d}{\epsilon^2}\log(\frac{d}{\epsilon}))$ by Theorems 5 and 4. □

Note, if a quadratic local model is constructed with $\mathcal{O}(d^2)$ order of interpolation points, then the oracle complexity increases to $w_\epsilon \leq \mathcal{O}(\frac{d^2}{\epsilon^2}\log(\frac{d}{\epsilon}))$, which is much worse. Hence, ASTRO-DF with enhancements can achieve the first-order stationary point faster than the original ASTRO-DF.

3 ASTRO-DF ENHANCEMENTS AND COMPLEXITY ANALYSIS

In ASTRO-DF the function values are estimated at each $\boldsymbol{x} \in \mathbb{R}^d$ with a random sample size of the form (4) that determines the accuracy needed at $\boldsymbol{x}$ to enable a speedy convergence to the first-order stationary point. More precisely, for the random sequence of incumbent solutions generated by ASTRO-DF that we denote by $\{\boldsymbol{X}_k, k \in \mathbb{N}\}$, each estimated objective function value is denoted by $f(\boldsymbol{X}_k, N_k)$. A complication in the stochastic setting is that unlike its deterministic counterpart, the function value estimates changes in almost all iterations, i.e., $f(\boldsymbol{X}_k, N_k) \neq f(\boldsymbol{X}_{k+1}, N_{k+1})$ even when $\boldsymbol{X}_k = \boldsymbol{X}_{k+1}$ as a result of unsuccessful iterations due to the changed sample size. The main changes from the original ASTRO-DF algorithm, as developed by Shashaani, Hashemi, and Pasupathy (2018) are the following:

1. In the new ASTRO-DF version, the point selection is deterministic. The only source of stochasticity is that only function estimates are used at each visited point instead of the true function values to build the local model, which determines the next iterate.
2. The local model is constructed with $2d$ new points along the coordinate bases on the trust-region boundary in each iteration to obtain a diagonal Hessian model.
3. The deterministic sample size lower bound sequence, i.e., $\{\lambda_k, k \in \mathbb{N}\}$ now grows logarithmically instead of linearly with k. Consequently, slower growth in the minimum sample size saves more of the budget as the iterations proceed. Showing that the new rate of λ_k is adequate for convergence requires an additional assumption on the function value observations. They are sub-exponential and possess well-defined moment-generating functions.

Algorithm 1 ASTRO-DF Algorithm with Enhancements

Require: Initial guess $\boldsymbol{x}_0 \in \mathbb{R}^d$, initial and maximum trust-region radii $0 < \Delta_0 < \Delta_{\max}$, model "fitness" thresholds $0 \leq \eta_1 < \eta_2$, trust-region contraction and expansion constants $0 < \gamma_1 < 1 < \gamma_2$, sample size lower bound sequence $\{\lambda_k\}_{k\in\mathbb{N}}$ such that $(\log k)^{1+\epsilon_0} = \mathcal{O}(\lambda_k)$ for some $\epsilon_0 > 0$, outer and inner adaptive sampling constants $\kappa_{\text{oas}}, \kappa_{\text{ias}} > 0$, inner loop contraction factor $w \in (0,1)$, and gradient balance constants $0 < \beta < \mu$.

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

Model Construction in a Contraction Loop:

2: Initialize $j = 1$.

3: **repeat**

4: Let $\Delta_k^{(j)} = \Delta_k w^{j-1}$, $\boldsymbol{X}_k := \boldsymbol{Y}_{k,0}^{(j)}$, and the interpolation set $\mathcal{Y}_{cb}^{(j)} = \left(\boldsymbol{Y}_{k,i}^{(j)} \in \mathcal{B}(\boldsymbol{X}_k; \Delta_k^{(j)})\right)_{i=1}^{2d}$.

5: Estimate $f\left(\boldsymbol{Y}_{k,i}^{(j)}, N\left(\boldsymbol{Y}_{k,i}^{(j)}\right)\right)$ for all $i = 0, 1, 2, \ldots, 2d$, where

$$N\left(\boldsymbol{Y}_{k,i}^{(j)}\right) = \max\left\{\lambda_k, \min\left\{n : \frac{\hat{\sigma}_F\left(\boldsymbol{Y}_{k,i}^{(j)}, n\right)}{\sqrt{n}} \leq \frac{\kappa_{\text{ias}}(\Delta_k^{(j)})^2}{\sqrt{\lambda_k}}\right\}\right\}. \tag{4}$$

6: Construct the model with a diagonal Hessian $M_k^{(j)}(\boldsymbol{X}_k + \boldsymbol{s})$ via interpolation and set $j = j + 1$.

7: **until** $\Delta_k^{(j)} \leq \mu \|\nabla M_k^{(j)}(\boldsymbol{X}_k)\|$.

8: Set $\check{N}_k = N(\boldsymbol{Y}_{k,0})$, and $\check{\Delta}_k = \min\left\{\Delta_k, \max\left\{\beta\left\|\nabla M_k^{(j)}(\boldsymbol{X}_k)\right\|, \Delta_k^{(j)}\right\}\right\}$.

TR Subproblem:

9: Approximate the k-th step $\boldsymbol{S}_k = \operatorname{argmin}_{\|\boldsymbol{s}\| \leq \check{\Delta}_k} M_k^{(j)}(\boldsymbol{X}_k + \boldsymbol{s})$, and set $\tilde{\boldsymbol{X}}_{k+1} = \boldsymbol{X}_k + \boldsymbol{S}_k$.

10: Estimate $f(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1})$, where

$$\widetilde{N}_{k+1} = \max\left\{\lambda_k, \min\left\{n : \frac{\hat{\sigma}_F\left(\tilde{\boldsymbol{X}}_{k+1}, n\right)}{\sqrt{n}} \leq \frac{\kappa_{\text{oas}}\check{\Delta}_k^2}{\sqrt{\lambda_k}}\right\}\right\}.$$

Update:

11: Compute success ratio $\hat{\rho}_k = \frac{f(\boldsymbol{X}_k, \check{N}_k) - f(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1})}{M_k^{(j)}(\boldsymbol{X}_k) - M_k^{(j)}(\tilde{\boldsymbol{X}}_{k+1})}$, and set

$$(\boldsymbol{X}_{k+1}, N_{k+1}, \Delta_{k+1}) = \begin{cases} (\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}, \min\{\gamma_2\check{\Delta}_k, \Delta_{\max}\}), & \text{if } \hat{\rho}_k \geq \eta_2 \text{ [very successful iteration]} \\ (\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}, \check{\Delta}_k), & \text{if } \hat{\rho}_k \in [\eta_1, \eta_2) \text{ [successful iteration]} \\ (\tilde{\boldsymbol{X}}_k, \check{N}_k, \gamma_1\check{\Delta}_k), & \text{otherwise [unsuccessful iteration].} \end{cases}$$

12: **end for**

4. Rather than two types of iterations, namely, successful and unsuccessful, we now allow success (accepting the candidate solution as new incumbent) to occur with minimal improvement instead of sufficient improvement. We measure the improvement by the success ratio using threshold η_1; note $\eta_1 = 0$ implies any improvement in the estimated function values is a success. However, we only expand the trust-region if the improvement is more significant using threshold $\eta_2 > \eta_1$. In this case, we call the iteration very successful, hence producing three types of iterations.

Algorithm 1 lists the new version with the enhancements. In what follows, we first show that the new local model construction with deterministic point selection renders the full-linearity of the model's

required quality in every iteration. This requirement in full Hessian quadratic models is expensive and hence only ensured in selected iterations or when near a critical region in the search space (Conn et al. 2009a). We extend the existing results that guarantee almost sure convergence of the algorithm to the enhanced ASTRO-DF. More importantly, we focus on the efficiency of the stochastic trust-region algorithm in terms of the expected number of iterations to reach an ϵ-accuracy in the optimality gap. To that end, we define the sample-path stopping time quantity $T_\epsilon := \min\{k : \|\nabla f(\boldsymbol{X}_k)\| \leq \epsilon\}$. Existing iteration complexity analysis by Blanchet et al. (2019) leverages a supermartingale property of the sample path optimality gap and provides $\mathcal{O}(\epsilon^{-2})$ iteration complexity in expectation. This result is proven for STORM and developed by Chen et al. (2018), that uses probabilistically accurate estimated and probabilistically fully-linear models. We extend the same results for the ASTRO-DF framework by proving the optimality gap stochastic process follows the supermartingale property-like conditions. However, we further discuss the requirements needed to make stronger iteration complexity with almost sure guarantees rather than in expectation guarantees, using similar arguments as in stochastic approximation and the famous result of Robbins and Siegmund (1971).

A chief assumption in obtaining the expected iteration complexity result for the STORM algorithm is that the estimated values are sufficiently accurate with a high probability. Our framework in ASTRO-DF justifies this assumption with probability one via our adaptive sampling rules. See Theorem 7 and remark afterward. Note, in the remainder of this paper, the estimation error is denote by $\bar{E}_k(N_k)$, i.e., $\bar{E}_k(N_k) = f(\boldsymbol{X}_k, N_k) - f(\boldsymbol{X}_k)$. The estimation error conditioned on the history will satisfy the following assumption:

Assumption 3 The Monte Carlo oracle generates iid random variables $F(\boldsymbol{X}_k, \xi_j) = f(\boldsymbol{X}_k) + E_j|\mathcal{F}_k$ with $\mathbb{E}[E_j|\mathcal{F}_k] = 0$, $\mathbb{E}[E_j^2|\mathcal{F}_k] = \sigma^2 < \infty$ for all k. Furthermore, the conditional errors $E_j|\mathcal{F}_k,\ j \geq 1$ are sub-exponential random variables almost surely, i.e., $\mathbb{E}[e^{tE_j}|\mathcal{F}_k] \leq e^{t^2K^2}$, for all t that satisfy $|t| \leq K^{-1}$ and some positive constant K.

As noted earlier, the estimation error is the only source of randomness in the new ASTRO-DF framework. From a similar analysis to Lemma 5.1 in (Shashaani et al. 2018) we obtain the desired result that states the estimation error rate of decay is eventually faster than Δ_k^2 with probability one:

Theorem 7 Let Assumptions 1 and 3 hold. Then for a given $c > 0$, we have that $|\bar{E}_k(N_k)| \leq c\check{\Delta}_k^2$ for sufficiently large k almost surely. In other words,

$$\Pr\{|\bar{E}_k(N_k)| \geq c\check{\Delta}_k^2 \text{ i.o.}\} = 0.$$

We refer the reader to (Vasquez, Shashaani, and Pasupathy 2021) for the detailed proofs. As a remark, note that if we select c large enough, then $-c\check{\Delta}_k^2 \leq \bar{E}_k(N_k) \leq c\check{\Delta}_k^2$ for all $k \in \mathbb{N}$ with probability one.

Let us now define a stochastic process $\{\tau_k, k \in \mathbb{N}\}$ that is formed on the basis of sample path optimality gap and trust-region size. In particular, we let $\tau_k := v(f(\boldsymbol{X}_k, N_k) - f^*) + (1-v)\Delta_k^2$ for some $v \in (0,1)$ and prove that it possesses a supermartingale property.

Theorem 8 Let Assumptions 1-3 hold and $\epsilon > 0$ be given. Then the sequence of iterates $\{\boldsymbol{X}_k, k \in \mathbb{N}\}$ generated by Algorithm 1 satisfies $\mathbb{E}[\tau_{k+1}|\mathcal{F}_k] \leq (\tau_k - \theta\Delta_k^2)$ for some constant $\theta > 0$. In other words, the expectation of the reduction in τ_k conditioned on the history is at most $-\theta\Delta_k^2$.

Proof. First, we note that

$$\mathbb{E}[\tau_{k+1} - \tau_k|\mathcal{F}_k] = v(\mathbb{E}[f(\boldsymbol{X}_{k+1}, N_{k+1}) - f(\boldsymbol{X}_k, N_k)|\mathcal{F}_k]) + (1-v)(\mathbb{E}[\Delta_{k+1}^2 - \Delta_k^2|\mathcal{F}_k]). \quad (5)$$

When k is unsuccessful, even though $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k$, since their adaptive sample sizes may vary, we can have that $f(\boldsymbol{X}_{k+1}, N_{k+1}) \neq f(\boldsymbol{X}_k, N_k)$. In this case, there exists a non-negative real number c (as defined in Theorem 7), such that the right hand side of (5) simplifies to $-(v - cv - 1)(\gamma_2^2 - 1)\Delta_k^2$. In the latter we have used the fact that $\check{\Delta}_{k+1} \leq \Delta_{k+1} \leq \gamma_2\check{\Delta}_k \leq \gamma_2\Delta_k$. When k is at least successful, i.e., $\hat{\rho}_k \geq \eta_1$,

the counterpart of Assumption 2 implies that

$$\begin{aligned} f(\boldsymbol{X}_{k+1}, N_{k+1}) - f(\boldsymbol{X}_k, N_k) &\leq \eta_1 (M_k(\boldsymbol{X}_{k+1}) - M_k(\boldsymbol{X}_k)) \\ &\leq -\frac{\kappa_{fcd}\eta_1}{2} \frac{\Delta_k w^{\kappa_{clu}-1}}{\mu} \min\{(\mu\kappa_{\mathcal{D}})^{-1}\Delta_k w^{\kappa_{clu}-1}, \check{\Delta}_k\} \\ &\leq -\Delta_k^2 (\frac{\kappa_{fcd}\eta_1}{2\mu} w^{2(\kappa_{clu}-1)} \min\{(\mu\kappa_{\mathcal{D}})^{-1}, 1\}), \end{aligned}$$

with the observation that $\check{\Delta}_k \geq \Delta_k w^{\kappa_{clu}-1}$. Hence, the right hand side of (5) this time simplifies to $-(v\frac{\kappa_{fcd}\eta_1}{2\mu} w^{2(\kappa_{clu}-1)} \min\{(\mu\kappa_{\mathcal{D}})^{-1}, 1\} + (v-1)(\gamma_2^2 - 1))\Delta_k^2$. To complete the proof, choose v such that $\max\{(v - cv - 1)(\gamma_2^2 - 1), v\frac{\kappa_{fcd}\eta_1}{2\mu} w^{2(\kappa_{clu}-1)} \min\{(\mu\kappa_{\mathcal{D}})^{-1}, 1\} + (v-1)(\gamma_2^2 - 1)\} \geq 0$.

□

The next new result proves the proposed ASTRO-DF variant's in expectation and almost sure complexity.

Theorem 9 Let Assumptions 1-3 hold and $\epsilon > 0$ be given. Then the following hold:

(i) $\mathbb{E}[T_\epsilon] = \mathcal{O}(\epsilon^{-2})$.
(ii) $\Pr\{\lim_{\epsilon\to 0} T_\epsilon \epsilon^2 = c' > 0\} = 1$ for some $c' > 0$.

Proof. Part (i) is a direct application of Theorem 2 in (Blanchet et al. 2019) that follows Theorem 8, namely, $\mathbb{E}[\tau_{k+1}|\mathcal{F}_k]\mathbb{I}\{k < T_\epsilon\} \leq (\tau_k - \theta\Delta_k^2)\mathbb{I}\{k < T_\epsilon\}$ for a given $\epsilon > 0$. We now prove part (ii) by defining the random variable $Z_k = \frac{1}{k} - c'\|\nabla f(\boldsymbol{X}_k)\|^2$ following the result of Theorem 8; observe that $Z_k \geq 0$ almost surely for large enough k. We now prove part (ii) by defining the random variable $Z_k = \frac{1}{k} - c'\|\nabla f(\boldsymbol{X}_k)\|^2$ following the result of Theorem 8; observe that $Z_k \geq 0$ almost surely for large enough c'. We can write

$$\begin{aligned} \mathbb{E}[Z_{k+1}|\mathcal{F}_k] &= \frac{1}{k+1} - c'\mathbb{E}[\|\nabla f(\boldsymbol{X}_{k+1})\|^2|\mathcal{F}_k] \\ &\leq \frac{1}{k}\left(1 - \frac{1}{k+1}\right) - c'\|\nabla f(\boldsymbol{X}_k)\|^2 \left(1 - \frac{2\mathbb{E}[\|\nabla f(\boldsymbol{X}_k) - \nabla f(\boldsymbol{X}_{k+1})\||\mathcal{F}_k]}{\|\nabla f(\boldsymbol{X}_k)\|}\right) \\ &\leq \left(1 - \frac{1}{k+1}\right)\left(\frac{1}{k} - c'\|\nabla f(\boldsymbol{X}_k)\|^2\right) = \left(1 - \frac{1}{k+1}\right) Z_k. \end{aligned}$$

In the above, we use two observations: (a) $\|\nabla f(\boldsymbol{X}_k)\| \leq \|\nabla M_k(\boldsymbol{X}_k)\| + \|\nabla f(\boldsymbol{X}_k) - \nabla M_k(\boldsymbol{X}_k)\| \leq (\frac{1}{\beta} + \kappa_{eg})\check{\Delta}_k$ due to full linearity of the model, and (b) $\mathbb{E}[\|\nabla f(\boldsymbol{X}_k) - \nabla f(\boldsymbol{X}_{k+1})\||\mathcal{F}_k] \leq \kappa_{Lg}\check{\Delta}_k$ due to the Lipschitz continuity of the gradient. Lastly, note that $(\frac{1}{\beta} + \kappa_{eg}) \leq 2(k+1)\kappa_{Lg}$ for sufficiently large k. Next, we apply the Robbins and Siegmund (1971) Theorem with $A_k = \frac{1}{K+1}$ and $B_k = C_k = 0$ to conclude that $Z_k \to 0$ as $k \to \infty$ almost surely. Alternatively, $(T_\epsilon)^{-1} - c'\epsilon^2 \to 0$ as $\epsilon \to 0$ with probability one, which in return completes the proof of part (ii). □

4 NUMERICAL EXPERIENCE

Problem	Description	Dimension	Known Structure
FACLOC	Facility Location	4	None
AMBUSQ	Ambulance Bases in a Square	6	f Discontinuous
SAN	Stochastic Activity Network Duration	13	f Convex

Table 1: The list of problems and their characteristics.

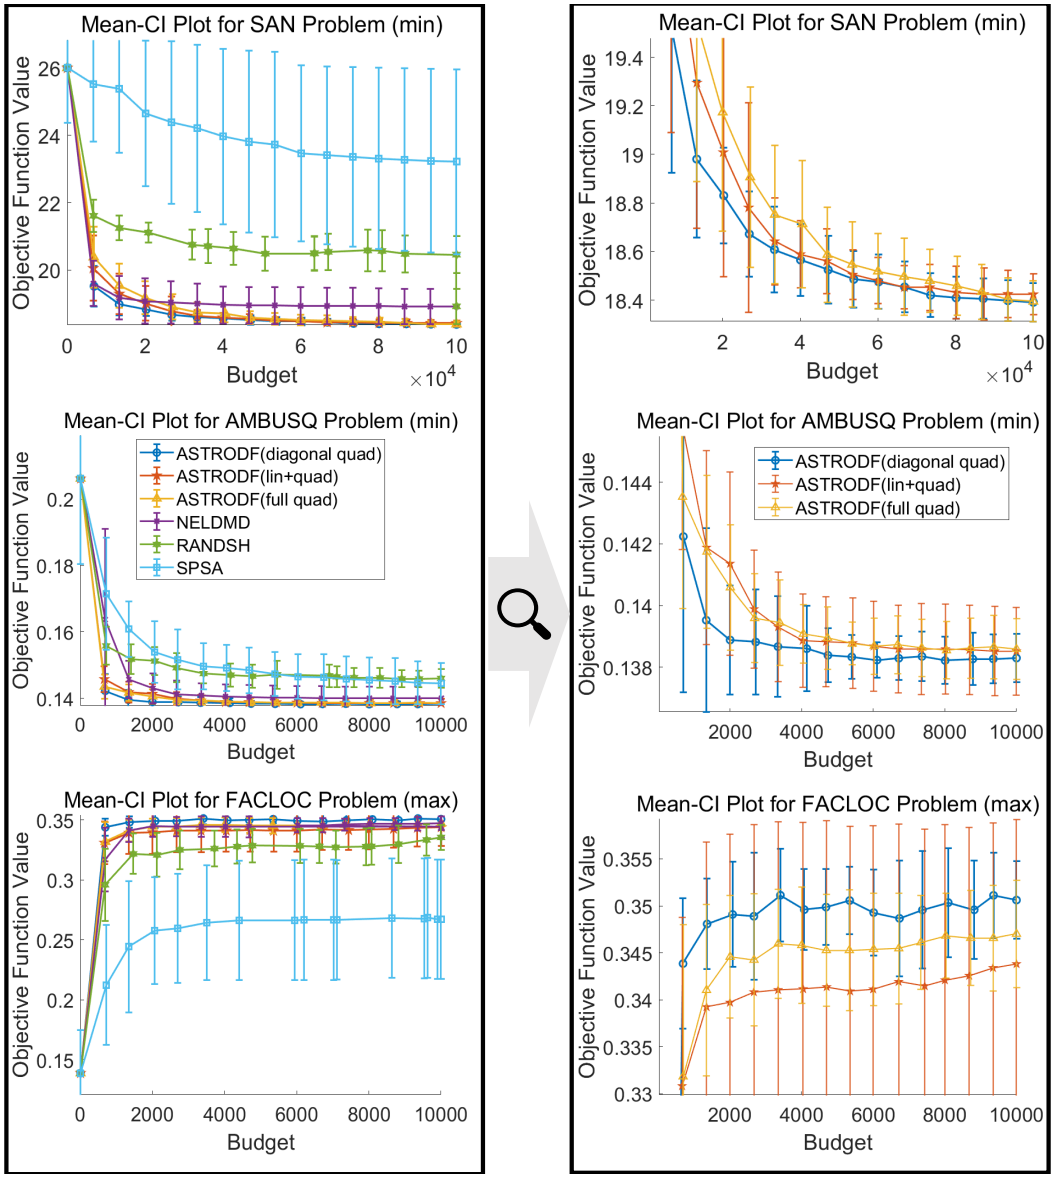

Figure 1: Performance of the 6 SO algorithms on 3 zeroth-order stochastic oracles of varying dimensions. At each budge point, an average of the estimated function value at each solution from 20 macro-replications is reported within its 95% confidence interval. The right panel is a zoomed-in view of the left.

In this section, we report the finite-time performance of the solvers on various problems from SimOpt library (Eckman et al. 2019). The SimOpt solver library contains direct-search methods, gradient-based methods, and model-based methods. The SimOpt problem library includes optimization problems in which the simulation oracle obtains objective function value at a given point. Hence, we can only know little information about the structure of the objective function for most problems in SimOpt. We would like to emphasize that utilizing stochastic simulation oracles in experimentation is preferred to the deterministic problems with added stochastic noise, for the latter leads to highly artificial solution-dependent estimators, especially with the use of CRN. Table 1 shows the SimOpt problems we use in this paper, all with unknown optimal solutions and varying simulation budgets. Since all algorithms work well on the low-dimensional problems we only report the higher dimensional problems in this paper.

There are three procedures in SimOpt. First, we run m number of macro-replications for each solver and problem via the `runwrapper`. Within a macro-replication, a solver solves the problem until the pre-defined budget is exhausted. At each solution $\boldsymbol{x}$, the function is evaluated by n number of replications to estimate the objective function with sample average approximation, which varies per solver (adaptive solvers use a random size $N(\boldsymbol{x})$). Second, the `postwrapper` runs ℓ post-replications at the intermediate solutions of each macro-replication for objective function estimates without optimization bias. Third, the `plotwrapper` records and plots the sample mean, sample variance, and quantiles at each of the solutions reported at the budget points across the macro-replications. SimOpt uses the budget points instead of the number of iterations since the running time depends on the number of oracle calls (i.e., the number of function evaluations). For a fair comparison and variance reduction, the SimOpt uses common random numbers for macro-replication and post-replications, however the initial solution changes at each macro-replication. We test the solvers' performance using the confidence intervals on each solver's produced solutions to understand their convergence behavior with $m = 20$ macro-replications and $\ell = 200$ post-replications.

For ASTRO-DF, we use following common parameters: $\mu = 100, \beta = 50, w = 0.9, \eta_1 = 0.1, \eta_2 = 0.5, \gamma_2 = 1.25^{(2/d)} = 1/\gamma_1, \lambda_k = 10(1 + (\log k)^{1.5}), \kappa_{\text{oas}} = \kappa_{\text{ias}} = 100$, and $\Delta_{\max} = 100$. For each macro-replication, we tune the Δ_0 by a pilot run for three candidates $0.08\Delta_{\max} \times (0.5^{\ln(d+1)}, 1, 1/0.5^{\ln(d+1)})$ using 1% of the total budget for each candidate. We compare two implementations of Algorithm 1 with that of the original ASTRO-DF - labeled by ASTRO-DF(full quad) - that always uses $(d+1)(d+2)/2$ poised points:

- ASTRO-DF(diagonal quad) always uses a quadraci local model with a diagonal Hessian to approximate *some* curvature information at a lower cost.
- ASTRO-DF(lin+quad) heuristically combines linear and fully quadratic models for the hope of improving practical efficiency. The heuristic ensures that local linear models are used when far away, i.e., $\Delta_k < \Delta_{\max}/100$ and quadratic models otherwise. We separately selected this heuristic condition among several others as the best.

For direct-search methods, we implemented the Nelder-Mead (NELDMD) and the Random Search (RANDSH). NELDMD iteratively maintains and updates a simplex of $d + 1$ vertices, and RANDSH evaluates solutions drawn from a particular probability distribution and is a global search method, whereas ASTRO-DF is a local search method. For gradient-based methods, we used the Simultaneous Perturbation Stochastic Approximation (SPSA). SPSA uses a line search method with the gradient estimation by the function values at two points: one in a random direction and another in the negative direction.

Figure 1 shows that among ASTRO-DF algorithms, ASTRO-DF(diagonal quad) has the fastest convergence rate and exhibits robust performance, i.e., achieves a good solution with high reliability (less variability = smaller CIs). NELDMD shows good performance in general. SPSA struggles relative to the other algorithms.

5 CONCLUSION

This paper leverages interpolation via coordinate bases in ASTRO-DF - an adaptive sampling solver based on the trust-region methodology for stochastic oracles. We show that the resulting underdetermined quadratic local model does not harm the consistency of ASTRO-DF. Our new variant substantially improves the computational complexity by reducing the cost of approximating the Hessian. Furthermore, we analyze the efficiency of the proposed algorithm and prove its almost sure $\mathcal{O}(\epsilon^{-2})$ iteration complexity. The original ASTRO-DF does not have any complexity guarantees. To the best of our knowledge, almost sure complexity results for this class of algorithms have not been obtained prior to this work. The interpretation of our new results is that while losing some of the model quality, applying the coordinate bases interpolation makes up for the high expense of handling the stochastic error by reducing the number of design points required from $\mathcal{O}(d^2)$ to $\mathcal{O}(d)$. Hence, ASTRO-DF with the new enhancements achieves the first-order stationary

point faster than the original ASTRO-DF. We leave the precise analysis of the oracle complexity to future research and discuss the new model construction approach's rough effect on the complexity results and our numerical experiments.

REFERENCES

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

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.

Conn, A. R., K. Scheinberg, and L. N. Vicente. 2009a. "Global convergence of general derivative-free trust-region algorithms to first-and second-order critical points". *SIAM Journal on Optimization* 20(1):387–415.

Conn, A. R., K. Scheinberg, and L. N. Vicente. 2009b. *Introduction to derivative-free optimization*. 1st ed. Society for Industrial and Applied Mathematics.

Coope, I. D., and R. Tappenden. 2020. "Gradient and Hessian approximations in derivative free optimization". *arXiv preprint arXiv:2001.08355*.

Curtis, F. E., and K. Scheinberg. 2020. "Adaptive Stochastic Optimization: A Framework for Analyzing Stochastic Optimization Algorithms". *IEEE Signal Processing Magazine* 37(5):32–42.

Eckman, D. J., S. G. Henderson, and R. Pasupathy. 2019. "Redesigning a testbed of simulation-optimization problems and solvers for experimental comparisons". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, , and Y.-J. Son, 3457–3467. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Robbins, H., and D. Siegmund. 1971. "A convergence theorem for non negative almost supermartingales and some applications". In *Optimizing methods in statistics*, 233–257. Elsevier.

Shashaani, S., F. 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.

Shashaani, S., S. R. Hunter, and R. Pasupathy. 2016. "ASTRO-DF: Adaptive Sampling Trust-Region Optimization Algorithms, Heuristics, and Numerical Experience.". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, and E. Zhou, 554–565. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

Vasquez, D., S. Shashaani, and R. Pasupathy. 2021. "The Complexity of Adaptive Sampling Trust-Region Methods for Nonconvex Stochastic Optimization". *Under Preparation*.

AUTHOR BIOGRAPHIES

YUNSOO HA is a Ph.D. student in the Department of Industrial and Systems Engineering at North Carolina State University. His research interests include simulation optimization, derivative-free optimization, and quantum computing. His e-mail address is yha3@ncsu.edu.

SARA SHASHAANI is an assistant professor in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. Her research interests are probabilistic data-driven modeling and simulation optimization. Her email address is sshasha2@ncsu.edu and her homepage is https://shashaani.wordpress.ncsu.edu/.

QUOC TRAN-DINH is an assistant professor at the Department of Statistics and Operations Research, The University of North Carolina at Chapel Hill. His research mainly focuses on numerical methods for continuous optimization, including convex, nonconvex, and stochastic optimization and related problems. His email address is quoctd@email.unc.edu.