

Expected complexity analysis of stochastic direct-search

Kwassi Joseph Dzahini*

March 6, 2020

Abstract: This work presents the convergence rate analysis of stochastic variants of the broad class of direct-search methods of directional type. It introduces an algorithm designed to optimize differentiable objective functions f whose values can only be computed through a stochastically noisy blackbox. The proposed stochastic directional direct-search (SDDS) algorithm accepts new iterates by imposing a sufficient decrease condition on so called probabilistic estimates of the corresponding unavailable objective function values. The accuracy of such estimates is required to hold with a sufficiently large but fixed probability β . The analysis of this method utilizes an existing supermartingale-based framework proposed for the convergence rates analysis of stochastic optimization methods that use adaptive step sizes. It aims to show that the expected number of iterations required to drive the norm of the gradient of f below a given threshold ϵ is bounded in $\mathcal{O}\left(\epsilon^{\frac{-p}{\min(p-1,1)}}/(2\beta-1)\right)$ with $p > 1$. Unlike prior analysis using the same aforementioned framework such as those of stochastic trust-region methods and stochastic line search methods, SDDS does not use any gradient information to find descent directions. However, its convergence rate is similar to those of both latter methods with a dependence on ϵ that also matches that of the broad class of deterministic directional direct-search methods which accept new iterates by imposing a sufficient decrease condition.

Keywords: Blackbox optimization, Derivative-free optimization, Stochastic optimization, Convergence rate, Direct-search, Stochastic processes.

*GERAD and Département de Mathématiques et de Génie Industriel, École Polytechnique de Montréal, C.P. 6079, Succ. Centre-ville, Montréal, Québec H3C 3A7, Canada (Kwassi-Joseph.Dzahini@polymtl.ca).

1 Introduction

Direct-search methods constitute a broad class of derivative-free optimization (DFO) methods where at each iteration, the DFO algorithm evaluates the objective function at a collection of points and acts solely based on those function values without any model building or derivative approximation [5, 10]. Such methods include as well those based on simplices like the classical Nelder-Mead method and its numerous variants, as those of directional type where an improvement in the objective function is guaranteed by moving along a direction defined by a better point [15].

This work focuses on the convergence rate analysis of stochastic variants of the broad class of directional direct-search methods analyzed in [15], using a supermartingale-based framework proposed in [8] and elements from [12]. It introduces a stochastic directional direct-search (SDDS) algorithm designed for stochastic blackbox optimization (BBO) and aims to solve the following unconstrained stochastic blackbox optimization problem which often arises in modern statistical machine learning:

$$\min_{x \in \mathbb{R}^n} f(x) \quad \text{with} \quad f(x) = \mathbb{E}_\Theta[f_\Theta(x)] \quad (1)$$

where Θ is a real-valued random variable following some unknown distribution, f_Θ denotes the black-box, the stochastically noisy computable version of the objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ which is numerically unavailable, and \mathbb{E}_Θ denotes the expectation with respect to Θ . Note that Θ is considered as a data point for many machine learning problems [14].

Significant theoretical and algorithmic advances have been made in the field of stochastic DFO in the recent years with the aim of solving Problem (1). Thus, numerous algorithms have been developed, most of which carry out either an estimation of the gradient of f using a single simulation, or a processing of the simulation model as a blackbox. However, since the simulation model can be inaccessible in many real applications, or the gradient can be too expensive to estimate computationally, direct-search optimization methods “appear to be the most promising option” [4].

Several recent works have proposed directional direct-search algorithms with full supported convergence rates analysis. Vicente [15] proved that to drive the gradient of an objective function below a threshold $\epsilon \in (0, 1)$, the number of iterations required by the broad class of directional direct-search methods that use a sufficient decrease condition when accepting new iterates, is bounded in $\mathcal{O}\left(\epsilon^{\frac{-p}{\min(p-1, 1)}}\right)$, with $p > 1$. Directional direct-search methods based on *probabilistic descent*, that incorporate random gradient, was recently proposed and analyzed by Gratton et al. [11] with worst-case complexity, and global rates results. However, both aforementioned previous works assume that the objective function is deterministic, i.e, function values are exactly computed.

Audet et al. [4] recently proposed StoMADS, a stochastic variant of the mesh adaptive direct-search (MADS) algorithm [3], with full-supported convergence analysis based on Clarke calculus and martingale theory. Using an algorithmic framework similar to that of MADS, Alarie et al. [1] also proposed another variant of MADS capable to optimize noisy blackboxes corrupted with Gaussian noise, and proved convergence results of the proposed method using statistical inference techniques. Nevertheless, no convergence rates analysis have been carried out for both methods.

The main novelty of the present work is that unlike many prior research on convergence rate analysis of stochastic DFO methods (see for example [6, 8, 14, 16] and references therein), especially those on stochastic trust-region [8] and line search [14] methods, SDDS does not use any first-order

information to find descent directions. Instead, such directions are provided by a positive spanning set and are chosen in such a way to ensure that they never become close to losing the positive spanning property. However, as emphasized in [15], “*it is not unreasonable*” to expect that SDDS shares a similar worst case complexity bound of the latter methods in term of the expected number of iterations. Indeed, one of the directions of any positive spanning set makes an acute angle with the negative gradient, provided that the objective function is continuously differentiable [12, 15]. This latter remark is in fact the cornerstone of the analysis in the present manuscript. Moreover, unlike the deterministic framework of [15], the proposed method accepts new iterates by imposing a sufficient decrease condition on so called *probabilistic estimates* of the corresponding unavailable objective function values, which accuracy is required to hold with a sufficiently large but fixed probability $\beta > 1/2$. However, even though such probability β of encountering sufficiently accurate estimates is not required to equal one, SDDS is shown to have desirable convergence properties. Specifically, as main theoretical result of the present work, the expected number of iterations required by SDDS to drive the gradient of f below a threshold ϵ is shown to be bounded in $\mathcal{O}\left(\epsilon^{\frac{1}{\min(p-1,1)}}/(2\beta - 1)\right)$, using a supermartingale-based framework proposed in [8]. Moreover, a subsequence of random iterates generated by SDDS is shown to drive the norm of the gradient of f to zero with probability one. Note also that the analysis is made very general in the present manuscript in the sense that it is not limited to $p = 2$, compared to several similar works, but instead, extends to $p > 1$. To the best of our knowledge, this research is the first to propose a convergence rate analysis of a stochastic direct-search algorithm of directional type.

This manuscript is organized as follows. Section 2 introduces an outline of the proposed stochastic algorithm and requirements on so-called probabilistic estimates that guarantee convergence at an appropriate rate. It is followed by Section 3 which presents a general framework of a stochastic process that is required to carry out the convergence rate analysis in Section 4. Section 4 also presents a lim inf-type first-order convergence result for SDDS, followed by a discussion and suggestions for future work.

2 The SDDS method and probabilistic estimates

SDDS, the stochastic algorithm analyzed in the present manuscript, is a direct-search method that uses inexact or noisy information about the objective f , specifically making use of so called probabilistic estimates. This section introduces the general framework of SDDS and discusses the requirements on the probabilistic estimates that guarantee the convergence of the algorithm.

2.1 The stochastic directional direct-search algorithm

The stochastic directional direct-search methods under study in the present manuscript use an algorithmic framework similar to that of the broad class of methods analyzed in [15], i.e., a framework that can describe the main features of generating set search (GSS) [12], pattern search and generalized pattern search (GPS) [2].

Each iteration of a directional direct-search method is composed of two main steps: the SEARCH step which is optional and the POLL step on which relies the convergence analysis. For simplicity

of presentation, Algorithm 1 does not show any SEARCH step. During the POLL, trial points are generated in a subset $\mathcal{P}^k = \{x^k + \delta^k d : d \in \mathbb{D}^k\}$ of the space of variables, where x^k denotes the incumbent solution, δ^k the step size and \mathbb{D}^k is a positive spanning set [5, 10]. Thus, the POLL step which follows stricter rules, consists of a local exploration of the variables space, unlike the SEARCH step which consists of a global exploration.

In Algorithm 1, since objective function values $f(x)$ are unavailable, f_0^k and f_s^k denote respectively the estimates of $f(x^k)$ and $f(x^k + s^k)$, with $s^k = \delta^k d$, constructed making use of evaluations of the noisy objective f_Θ . In order for the information provided by f_0^k and f_s^k to determine the iteration type, i.e., successful or unsuccessful, both estimates are required to be ε_f -accurate, with $\varepsilon_f > 0$, according to the following definition similar to those in [4, 8, 9, 14].

Definition 1. Let $\rho : (0, +\infty) \rightarrow (0, +\infty)$ be a continuous and non-decreasing function satisfying $\rho(t)/t \rightarrow 0$ when $t \searrow 0$. f^k is called ε_f -accurate estimate of $f(x^k)$ for a given δ^k if

$$|f^k - f(x^k)| \leq \varepsilon_f \rho(\delta^k).$$

Following the terminology in [12], the function ρ in Definition 1 represents the “forcing function”. Sufficient information to determine the iteration type is provided next.

Proposition 1. Let f_0^k and f_s^k be ε_f -accurate estimates of $f(x^k)$ and $f(x^k + s^k)$ respectively, and let $\gamma > 2$ be a fixed constant. Then the followings hold:

$$\text{if } f_s^k - f_0^k \leq -\gamma \varepsilon_f \rho(\delta^k), \text{ then } f(x^k + s^k) - f(x^k) \leq -(\gamma - 2) \varepsilon_f \rho(\delta^k) := u_{\text{sto}}^k \quad (2)$$

$$\text{if } f_s^k - f_0^k > -\gamma \varepsilon_f \rho(\delta^k), \text{ then } f(x^k + s^k) - f(x^k) > -(\gamma + 2) \varepsilon_f \rho(\delta^k) := \ell_{\text{sto}}^k \quad (3)$$

Proof. The proof straightforwardly follows from Definition 1 and the equality

$$f(x^k + s^k) - f(x^k) = f(x^k + s^k) - f_s^k + (f_s^k - f_0^k) + f_0^k - f(x^k).$$

□

In addition to the results in Proposition 1, the definition of the iteration type in Algorithm 1 is motivated by the following remarks. First, notice that in the stochastic framework of StomADS where $\rho(t) = t^2$, since as in (3) the inequality $f_s^k - f_0^k > -\gamma \varepsilon_f \rho(\delta_p^k)$ (δ_p^k denoting the so-called frame size parameter) does not necessarily lead to an increase in the unavailable objective function f , two types of unsuccessful iterations have been distinguished. Unsuccessful iterations which are called *certain*, are characterized by $f_s^k - f_0^k \geq -\gamma \varepsilon_f \rho(\delta_p^k)$ and lead to an increase in f whenever both estimates f_0^k and f_s^k are accurate, while those such that $-\gamma \varepsilon_f \rho(\delta_p^k) < f_s^k - f_0^k < \gamma \varepsilon_f \rho(\delta_p^k)$ are called *uncertain* since they lead to $-(\gamma + 2) \varepsilon_f \rho(\delta^k) < f(x^k + s^k) - f(x^k) < (\gamma + 2) \varepsilon_f \rho(\delta^k)$. Then, even though updating the frame size parameter according to $\delta_p^{k+1} = \tau \delta_p^k$ on *uncertain* unsuccessful iterations, and $\delta_p^{k+1} = \tau^2 \delta_p^k$ whenever the unsuccessful iteration is *certain* ($\tau \in (0, 1)$ being a rational number), the corresponding sequence $\{\delta_p^k\}_{k \in \mathbb{N}}$ was shown in [4] to converge to zero. Note also that this kind of update is the only one that differentiates *certain* iterations from those that are *uncertain*. In the present work, the step size parameter δ^k is therefore updated on unsuccessful iterations according to $\delta^{k+1} = \tau \delta^k$, where τ is a real number in $(0, 1)$. As a consequence, *certain* unsuccessful iterations will not be differentiated

from *uncertain* ones. In other words, every iteration such that $f_s^k - f_0^k > -\gamma\varepsilon_f\rho(\delta^k)$ will be called unsuccessful.

However, let put an emphasis on the specific choice of τ by means of the following additional remarks. Note that in the general deterministic framework described in [15], the amount of decrease in the objective function on successful iterations is such that $f(x^k + s^k) - f(x^k) \leq -\rho(\delta^k) := u_{\text{det}}^k$ while unsuccessful iterations are characterized by $f(x^k + s^k) - f(x^k) > -\rho(\delta^k) := \ell_{\text{det}}^k$. Thus, the equality $\ell_{\text{det}}^k = u_{\text{det}}^k$ always holds, which is not the case in stochastic settings where $\ell_{\text{sto}}^k < u_{\text{sto}}^k$. Moreover, since $\delta^{k+1} < \delta^k$ whenever the iteration k is unsuccessful, then $\ell_{\text{det}}^{k+1} > u_{\text{det}}^k$. Likewise, since $\delta^{k+1} > \delta^k$ on successful iterations, then $u_{\text{det}}^{k+1} < \ell_{\text{det}}^k$. Given that the equality $\ell_{\text{sto}}^k = u_{\text{sto}}^k$ can not hold in the present stochastic settings, then τ must be chosen in such a way that at least, both inequalities $\ell_{\text{sto}}^{k+1} > u_{\text{sto}}^k$ and $u_{\text{sto}}^{k+1} < \ell_{\text{sto}}^k$ hold respectively on unsuccessful and successful iterations, analogously to the deterministic framework. This means using (2) and (3), that τ must be chosen according to

$$\rho(\tau\delta^k) < \frac{\gamma - 2}{\gamma + 2}\rho(\delta^k) \quad \text{and} \quad \rho(\tau^{-1}\delta^k) > \frac{\gamma + 2}{\gamma - 2}\rho(\delta^k). \quad (4)$$

It follows from (4) that depending on the expression of the forcing function ρ , the choice of τ could depend on δ^k and hence should be made at each iteration. Thus, in order to make the present analysis simpler, the following assumption is made.

Assumption 1. *The forcing function $\rho: (0, +\infty) \rightarrow (0, +\infty)$ is such that $\rho(t) = ct^p$, where $c > 0$ and $p > 1$ are fixed constants.*

Under Assumption 1, the choice of τ does not depend on δ^k . More precisely, it follows from (4) that τ must be chosen according to $0 < \tau^p < \frac{\gamma-2}{\gamma+2}$, for all $k \in \mathbb{N}$, as specified in Algorithm 1.

2.2 Probabilistic estimates

Following the notation in [7], all stochastic quantities in the present manuscript live on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a nonempty set referred to as the sample space, \mathcal{F} is a collection of events (subsets of Ω) called a σ -field and \mathbb{P} is a finite measure on the measurable space (Ω, \mathcal{F}) satisfying $\mathbb{P}(\Omega) = 1$ and referred to as probability measure. The elements $\omega \in \Omega$ are referred to as possible outcomes or sample points. When \mathbb{R}^n is given its Borel σ -field, i.e., the one generated by the open sets, a random variable or random map X is a measurable map on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ into the measurable space $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Measurability meaning that each event $\{X \in I\} := X^{-1}(I)$ belongs to \mathcal{F} for all $I \in \mathcal{B}(\mathbb{R}^n)$ [7].

The estimates f_s^k and f_0^k constructed at iteration k of Algorithm 1, based on random information provided by the noisy objective f_Θ , can be considered as realizations of random estimates F_s^k and F_0^k respectively. Thus, because of the randomness stemming from such random estimates whose behavior influences each iteration k , Algorithm 1 results in a stochastic process $\{X^k, S^k, \Delta^k, F_s^k, F_0^k\}$. In general, uppercase letters will be used to denote random variables while lowercase letters will be used for their realizations. For example, $x^k = X^k(\omega)$, $s^k = S^k(\omega)$ and $\delta^k = \Delta^k(\omega)$ denote respectively realizations of the random variables X^k , S^k and Δ^k . Similarly, following the notations in [4, 8, 9, 14], $f_0^k = F_0^k(\omega)$ and $f_s^k = F_s^k(\omega)$ with F_0^k and F_s^k denoting respectively estimates of $f(X^k)$ and $f(X^k + S^k)$.

Algorithm 1: SDDS

1 [0] Initialization

2 Choose $x^0 \in \mathbb{R}^n$, $\delta^0 > 0$, $\varepsilon_f > 0$, $\gamma > 2$, $c > 0$, $p > 1$, $0 < \tau < \left(\frac{\gamma-2}{\gamma+2}\right)^{1/p}$, $j_{\max} \in \mathbb{N}$

3 and $\delta_{\max} = \tau^{-j_{\max}} \delta^0$.

4 Set the iteration counter $k \leftarrow 0$.

5 [1] Poll

6 Select a positive spanning set \mathbb{D}^k .

7 Generate a set \mathcal{P}^k of Poll points such that $\mathcal{P}^k = \{x^k + \delta^k d : d \in \mathbb{D}^k\}$.

8 Obtain estimates f_0^k and f_s^k of $f(x^k)$ and $f(x^k + s^k)$, respectively, using objective function evaluations.

10 Success

11 If $f_s^k - f_0^k \leq -\gamma c \varepsilon_f (\delta^k)^p$ for some $s^k = \delta^k d^k \in \{\delta^k d : d \in \mathbb{D}^k\}$,

12 Set $x^{k+1} \leftarrow x^k + s^k$, and $\delta^{k+1} \leftarrow \min\{\tau^{-1} \delta^k, \delta_{\max}\}$.

13 Failure

14 Otherwise set $x^{k+1} \leftarrow x^k$ and $\delta^{k+1} \leftarrow \tau \delta^k$.

15 [2] Termination

16 If no termination criterion is met,

17 Set $k \leftarrow k + 1$ and go to **[1]**.

18 Otherwise stop.

Figure 1: Pseudo code of the Stochastic Directional Direct-Search (SDDS) algorithm. Success or failure is determined during the Poll at iteration k , using information provided by both estimates f_0^k and f_s^k in order to update the step size parameter δ^k and the current iterate x^k . As long as no stopping criterion is met, a new iteration is initiated with a new step size parameter δ^{k+1} .

The goal of this work is to show that the stochastic process resulting from Algorithm 1 converges at an appropriate rate with probability one, provided that the sequence $\{(F_0^k, F_s^k)\}$ is sufficiently accurate with sufficiently high but fixed probability, conditioned on the past.

As proposed in [9, 14], the notion of conditioning on the past is formalized in the following definition similar to those in [4, 8, 9, 14, 16], where \mathcal{F}_{k-1}^F denotes the σ -field generated by $F_0^0, F_s^0, F_0^1, F_s^1, \dots, F_0^{k-1}$ and F_s^{k-1} , with \mathcal{F}_{-1}^F being set to equal $\sigma(x^0)$ for completeness. Thus, one can notice that $\mathbb{E}(\Delta^k | \mathcal{F}_{k-1}^F) = \Delta^k$ and $\mathbb{E}(X^k | \mathcal{F}_{k-1}^F) = X^k$ for all $k \geq 0$, by construction of the random variables Δ^k and X^k in Algorithm 1.

Definition 2. A sequence of random estimates $\{(F_0^k, F_s^k)\}$ is said to be β -probabilistically ε_f -accurate with respect to the corresponding sequence $\{X^k, S^k, \Delta^k\}$ if the events

$$J_k = \{F_0^k, F_s^k, \text{ are } \varepsilon_f\text{-accurate estimates of } f(x^k) \text{ and } f(x^k + s^k), \text{ respectively}\}$$

satisfy the following submartingale-like condition

$$\mathbb{P}(J_k | \mathcal{F}_{k-1}^F) = \mathbb{E}(\mathbb{1}_{J_k} | \mathcal{F}_{k-1}^F) \geq \beta,$$

where $\mathbb{1}_{J_k}$ denotes the indicator function of the event J_k , that is $\mathbb{1}_{J_k} = 1$ if $\omega \in J_k$ and 0 otherwise.

An estimate is called “good” if $\mathbb{1}_{J_k} = 1$. Otherwise it is called “bad” [4].

Global convergence properties of deterministic directional direct-search methods strongly rely on having the step size parameters approaching zero [15] and the fact that the function value $f(x)$ never increases after an iteration. The main challenge of the analysis in the present stochastic framework lies in the fact that this monotonicity is not always guaranteed. The key to the analysis of Algorithm 1 thus relies on the assumption that accuracy in function estimates “improves in coordination with the perceived progress of the algorithm” [8]. The analysis is based on properties of supermartingales whose increments have a decreasing tendency and depend on the change in objective function values between iterations.

In order to show that the sequence $\{\Delta^k\}_{k \in \mathbb{N}}$ of random step size parameters converges to zero with probability one, let make the following key assumption similar to those in [4, 14].

Assumption 2. For some fixed $\beta \in (0, 1)$, and $\varepsilon_f > 0$, the followings hold for the random quantities derived from Algorithm 1.

(i) The sequence $\{(F_0^k, F_s^k)\}$ of estimates is β -probabilistically ε_f -accurate.

(ii) The sequence $\{(F_0^k, F_s^k)\}$ satisfies the following variance condition

$$\begin{aligned} \mathbb{E} \left(|F_0^k - f(X^k)|^2 \mid \mathcal{F}_{k-1}^F \right) &\leq \varepsilon_f^2 (1 - \beta) [\rho(\Delta^k)]^2 \\ \text{and} \quad \mathbb{E} \left(|F_s^k - f(X^k + S^k)|^2 \mid \mathcal{F}_{k-1}^F \right) &\leq \varepsilon_f^2 (1 - \beta) [\rho(\Delta^k)]^2 \end{aligned} \quad (5)$$

By means of Assumption 2-(ii), the variance in function estimates is adaptively controlled. Showing therefore that the sequence of random step size parameters converges to zero with probability one, ensures that this variance is driven to zero even though the probability β of encountering good estimates remains fixed, thus allowing Algorithm 1 to behave like an exact deterministic method asymptotically.

Moreover, since the estimates satisfying Assumption 2 can easily be constructed using techniques proposed in [8, 9, 14], then thorough details about their computations are not provided here again. Note however that if Θ^0 and Θ^s are two independent random variables following the same distribution as Θ defined in (1), and if $\Theta_i^0, i = 1, 2, \dots, p^k$ and $\Theta_j^s, j = 1, 2, \dots, p^k$ are independent random samples of Θ^0 and Θ^s respectively, then the estimates

$$F_0^k = \frac{1}{p^k} \sum_{i=1}^{p^k} f_{\Theta_i^0}(x^k) \quad \text{and} \quad F_s^k = \frac{1}{p^k} \sum_{j=1}^{p^k} f_{\Theta_j^s}(x^k + s^k)$$

satisfy Assumption 2 provided that the sample size p^k satisfies

$$p^k \geq \frac{V}{\varepsilon_f^2 (1 - \sqrt{\beta}) [\rho(\delta^k)]^2},$$

where the constant $V > 0$ is such that the variance of $f_{\Theta}(x)$ satisfies $\mathbb{V}[f_{\Theta}(x)] \leq V < +\infty$, for all $x \in \mathbb{R}^n$.

Next is stated a useful lemma similar to those in [4, 14], linking the probability of obtaining bad estimates to the variance assumption on function values.

Lemma 1. *Let Assumption 2 holds. Then for all $k \geq 0$, the followings hold for the random process $\{X^k, F_0^k, F_s^k, \Delta^k\}$ generated by Algorithm 1*

$$\begin{aligned} & \mathbb{E} \left(\mathbb{1}_{\bar{J}_k} |F_0^k - f(X^k)| \mid \mathcal{F}_{k-1}^F \right) \leq \varepsilon_f (1 - \beta) [\rho(\Delta^k)] \\ \text{and} \quad & \mathbb{E} \left(\mathbb{1}_{\bar{J}_k} |F_s^k - f(X^k + S^k)|^2 \mid \mathcal{F}_{k-1}^F \right) \leq \varepsilon_f (1 - \beta) [\rho(\Delta^k)] \end{aligned} \quad (6)$$

Proof. The result is proved using ideas derived from [4, 14]. The proof follows straightforwardly from the conditional Cauchy-Schwarz inequality [7] as follows

$$\begin{aligned} \mathbb{E} \left(\mathbb{1}_{\bar{J}_k} |F_s^k - f(X^k + S^k)| \mid \mathcal{F}_{k-1}^F \right) & \leq \left[\mathbb{E} \left(\mathbb{1}_{\bar{J}_k} \mid \mathcal{F}_{k-1}^F \right) \right]^{1/2} \left[\mathbb{E} \left(|F_s^k - f(X^k + S^k)|^2 \mid \mathcal{F}_{k-1}^F \right) \right]^{1/2} \\ & \leq (1 - \beta)^{1/2} \varepsilon_f (1 - \beta)^{1/2} [\rho(\Delta^k)], \end{aligned}$$

where the last inequality follows from (5) and the fact that $\mathbb{E} \left(\mathbb{1}_{\bar{J}_k} \mid \mathcal{F}_{k-1}^F \right) = \mathbb{P} \left(\mathbb{1}_{\bar{J}_k} \mid \mathcal{F}_{k-1}^F \right) \leq 1 - \beta$ thanks to Assumption 2-(i). The proof for $F_0^k - f(X^k)$ is the same. \square

3 A renewal-reward martingale process

This section presents a general stochastic process and its associated stopping time T introduced in [8] for the convergence rate analysis of a stochastic trust-region method. It introduces some relevant definition, assumptions and theorem derived in the analysis of a renewal-reward process in [8], that will be useful for the convergence rate analysis presented in Section 4. Specifically, by considering the stopping time consisting of the time required by SDDS to reach a desired accuracy, Section 4 will aim to show how the properties of this general stochastic process are satisfied for Algorithm 1. Note that some results derived in analyzing this stochastic process in [8] are also used in [14] for the convergence rate analysis of a stochastic line search method.

Definition 3. *A random variable T is said to be a stopping time with respect to a given discrete time stochastic process $\{X_k\}_{k \in \mathbb{N}}$ if, for each $k \in \mathbb{N}$, the event $\{T = k\}$ belongs to the σ -field $\sigma(X_1, X_2, \dots, X_k)$ generated by X_1, X_2, \dots, X_k .*

Consider a stochastic process $\{(\Phi_k, \Delta^k)\}_{k \in \mathbb{N}}$ satisfying $\Phi_k \in [0, +\infty)$ and $\Delta^k \in [0, +\infty)$ for all $k \in \mathbb{N}$. Define on the same probability space as $\{(\Phi_k, \Delta^k)\}_{k \in \mathbb{N}}$, a sequence of biased random walk process $\{W_k\}_{k \in \mathbb{N}}$ such that $W_0 = 1$,

$$\mathbb{P}(W_{k+1} = 1 \mid \mathcal{F}_k) = q \quad \text{and} \quad \mathbb{P}(W_{k+1} = -1 \mid \mathcal{F}_k) = 1 - q, \quad (7)$$

where $q \in (1/2, 1)$ and \mathcal{F}_k denotes the σ -field generated by $\{(\Phi_0, \Delta^0, W_0), (\Phi_1, \Delta^1, W_1), \dots, (\Phi_k, \Delta^k, W_k)\}$.

Define the following family $\{T_{\epsilon'}\}_{\epsilon' > 0}$ of stopping times parameterized by $\epsilon' > 0$, with respect to $\{\mathcal{F}_k\}_{k \in \mathbb{N}}$. The following assumptions are made in [8, 14] in order to derive a bound on $\mathbb{E}(T_{\epsilon'})$.

Assumption 3. *The following hold for the stochastic process $\{(\Phi_k, \Delta^k, W_k)\}_{k \in \mathbb{N}}$.*

- (i) *There exist constant $\lambda \in (0, +\infty)$ and $\delta_{\max} = \delta^0 e^{\lambda j_{\max}}$, for some integer $j_{\max} \in \mathbb{Z}$, such that $\Delta^k \leq \delta_{\max}$ for all $k \in \mathbb{N}$.*

(ii) There exists a constant $\delta_{\epsilon'} = \delta^0 e^{\lambda j_{\epsilon'}}$, for some $j_{\epsilon'} \in \mathbb{Z}$, $j_{\epsilon'} \leq 0$, such that the following holds for all $k \in \mathbb{N}$,

$$\mathbb{1}_{\{T_{\epsilon'} > k\}} \Delta^{k+1} \geq \mathbb{1}_{\{T_{\epsilon'} > k\}} \min(\Delta^k e^{\lambda W_{k+1}}, \delta_{\epsilon'}), \quad (8)$$

where W_{k+1} satisfies (7) with $q > \frac{1}{2}$.

(iii) There exists a nondecreasing function $h : [0, +\infty) \rightarrow (0, +\infty)$ and a constant $\eta > 0$ such that

$$\mathbb{E}(\Phi_{k+1} - \Phi_k \mid \mathcal{F}_k) \mathbb{1}_{\{T_{\epsilon'} > k\}} \leq -\eta h(\Delta^k) \mathbb{1}_{\{T_{\epsilon'} > k\}}. \quad (9)$$

Note that as highlighted in [8, 14], Assumption 3 states that conditioned on the past, the nonnegative random sequence $\{\Phi_k\}_{k \in \mathbb{N}}$ decreases by at least $\eta h(\Delta^k)$ at each iteration provided that $T_{\epsilon'} > k$ and moreover, the sequence $\{\Delta^k\}_{k \in \mathbb{N}}$ has a tendency to increase whenever it is below some fixed threshold $\delta_{\epsilon'}$.

The following theorem providing a bound on $\mathbb{E}(T_{\epsilon'})$ is proved in [8] by observing that the upward drift in the random walk $\{W_k\}_{k \in \mathbb{N}}$ makes the event $\{\Delta^k \geq \delta_{\epsilon'}\}$ occur sufficiently frequently on average [8, 14]. Hence, $\mathbb{E}(\Phi_{k+1} - \Phi_k)$ can frequently be bounded by some negative fixed constant, thus leading to a bound on the expected stopping time $\mathbb{E}(T_{\epsilon'})$.

Theorem 1. *Let Assumption 3 hold. Then,*

$$\mathbb{E}(T_{\epsilon'}) \leq \frac{q}{2q-1} \times \frac{\Phi_0}{\eta h(\delta_{\epsilon'})} + 1$$

4 Convergence rate analysis

It follows from Section 3 that Theorem 1 holds for any stopping time $T_{\epsilon'}$ defined with respect to the filtration $\{\mathcal{F}_k\}_{k \in \mathbb{N}}$, provided that Assumption 3 hold for the stochastic process $\{(\Phi_k, \Delta^k, W_k)\}_{k \in \mathbb{N}}$. Thus, the goal of the present section is to show how such a stochastic process satisfying Assumption 3 can be constructed in order to bound the expected number of iterations required by Algorithm 1 to achieve $\|\nabla f(X^k)\| \leq \epsilon$, for some arbitrary fixed $\epsilon \in (0, 1)$, where $\|\cdot\|$ denotes the Euclidean norm of \mathbb{R}^n as in the remainder of the manuscript.

4.1 Analysis of the stochastic process generated by SDDS

In order to show that Assumption 3 holds, let impose the following standard assumption on the objective function f .

Assumption 4. *The function f is bounded from below, i.e., there exists $f_{\min} \in \mathbb{R}$ such that $-\infty < f_{\min} \leq f(x)$, for all $x \in \mathbb{R}^n$.*

The following result generalizing that in [4] provides a bound on the expected decrease in the random function

$$\Phi_k := \frac{\nu}{c\mathcal{E}_f} (f(X^k) - f_{\min}) + (1 - \nu)(\Delta^k)^p. \quad (10)$$

Theorem 2. Let Assumption 1, 2 and 4 hold. Let $\gamma > 2, p > 1$ and $\tau \in (0, 1)$. Let $\nu \in (0, 1)$ and $\beta \in (1/2, 1)$ be chosen such that

$$\frac{\nu}{1-\nu} \geq \frac{2(\tau^{-p}-1)}{\gamma-2} \quad \text{and} \quad \frac{\beta}{1-\beta} \geq \frac{\nu}{1-\nu} \times \frac{4}{(1-\tau^p)}, \quad (11)$$

Then the expected decrease in the random function Φ_k defined in (10) satisfies

$$\mathbb{E}(\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1}^F) \leq -\frac{1}{2}\beta(1-\nu)(1-\tau^p)(\Delta^k)^p. \quad (12)$$

Proof. The proof is almost identical to that in [4], using ideas derived in [9, 13, 14] and making use of properties of the random function Φ_k defined in (10). It considers two separate cases: good estimates and bad estimates, each of which are broken into whether an iteration is successful or unsuccessful. Define the event S by

$$S := \{\text{The iteration is successful}\},$$

and let \bar{S} denote the complement of S .

Case 1 (Good estimates, $\mathbb{1}_{J_k} = 1$) The overall goal is to show that Φ_k decreases no matter what type of iteration occurs thus yielding the following bound

$$\mathbb{E}(\mathbb{1}_{J_k}(\Phi_{k+1} - \Phi_k) \mid \mathcal{F}_{k-1}^F) \leq -\beta(1-\nu)(1-\tau^p)(\Delta^k)^p. \quad (13)$$

(i) *Successful iteration* ($\mathbb{1}_S = 1$). A decrease occurs in f according to (2) since estimates are good and the iteration is successful, thus implying that

$$\mathbb{1}_{J_k} \mathbb{1}_S \frac{\nu}{c\varepsilon_f} (f(X^{k+1}) - f(X^k)) \leq -\mathbb{1}_{J_k} \mathbb{1}_S \nu(\gamma-2)(\Delta^k)^p. \quad (14)$$

The step size parameter is updated according to $\Delta^{k+1} = \min\{\tau^{-1}\Delta^k, \delta_{\max}\}$. Hence,

$$\mathbb{1}_{J_k} \mathbb{1}_S (1-\nu) [(\Delta^{k+1})^p - (\Delta^k)^p] \leq \mathbb{1}_{J_k} \mathbb{1}_S (1-\nu)(\tau^{-p}-1)(\Delta^k)^p. \quad (15)$$

Then, choosing ν according to (11) ensures that the right-hand side term of (14) dominates that of (15), i.e.,

$$-\nu(\gamma-2)(\Delta^k)^p + (1-\nu)(\tau^{-p}-1)(\Delta^k)^p \leq -\frac{1}{2}\nu(\gamma-2)(\Delta^k)^p. \quad (16)$$

Thus, combining (14), (15) and (16) yields

$$\mathbb{1}_{J_k} \mathbb{1}_S (\Phi_{k+1} - \Phi_k) \leq -\mathbb{1}_{J_k} \mathbb{1}_S \frac{1}{2}\nu(\gamma-2)(\Delta^k)^p. \quad (17)$$

(ii) *Unsuccessful iteration* ($\mathbb{1}_{\bar{S}} = 1$). The step size parameter is decreased while there is a change of zero in function values since the iteration is unsuccessful. Thus,

$$\mathbb{1}_{J_k} \mathbb{1}_{\bar{S}} (\Phi_{k+1} - \Phi_k) = -\mathbb{1}_{J_k} \mathbb{1}_{\bar{S}} (1-\nu)(1-\tau^p)(\Delta^k)^p. \quad (18)$$

Then, choosing ν according to (11) and noticing that $1 - \tau^p < \tau^{-p} - 1$, ensure that unsuccessful iterations, specifically (18), provide the worst case decrease when compared to (17), i.e., the following holds

$$-\frac{1}{2}\nu(\gamma - 2)(\Delta^k)^p \leq -(1 - \nu)(1 - \tau^p)(\Delta^k)^p. \quad (19)$$

Thus, combining (17), (18), and (19), leads to the following bound on the change in Φ_k

$$\mathbb{1}_{J_k}(\Phi_{k+1} - \Phi_k) = \mathbb{1}_{J_k}(\mathbb{1}_S + \mathbb{1}_{\bar{S}})(\Phi_{k+1} - \Phi_k) \leq -\mathbb{1}_{J_k}(1 - \nu)(1 - \tau^p)(\Delta^k)^p. \quad (20)$$

Since Assumption 2 holds, then taking conditional expectations with respect to \mathcal{F}_{k-1}^F in both sides of (20) leads to (13).

Case 2 (Bad estimates, $\mathbb{1}_{\bar{J}_k} = 1$). Since the estimates are bad, an iterate leading to an increase in f and Δ^k , and hence in Φ_k , can be accepted by Algorithm 1. Such an increase in Φ_k is controlled by bounding the variance in function estimates, using (5). Then, in order to guarantee that Φ_k is sufficiently reduced in expectation, the probability of outcome is adjusted to be sufficiently small. The overall goal is to show that

$$\mathbb{E}(\mathbb{1}_{\bar{J}_k}(\Phi_{k+1} - \Phi_k) \mid \mathcal{F}_{k-1}^F) \leq 2\nu(1 - \beta)(\Delta^k)^p. \quad (21)$$

(i) *Successful iteration* ($\mathbb{1}_S = 1$). The change in f is bounded as follows

$$\begin{aligned} \mathbb{1}_{\bar{J}_k} \mathbb{1}_S \frac{\nu}{c\varepsilon_f} (f(X^{k+1}) - f(X^k)) &\leq \mathbb{1}_{\bar{J}_k} \mathbb{1}_S \frac{\nu}{c\varepsilon_f} [(F_s^k - F_0^k) + |f(X^{k+1}) - F_s^k| + |F_0^k - f(X^k)|] \\ &\leq \mathbb{1}_{\bar{J}_k} \mathbb{1}_S \nu \left[-\gamma(\Delta^k)^p + \frac{1}{c\varepsilon_f} (|f(X^{k+1}) - F_s^k| + |F_0^k - f(X^k)|) \right] \end{aligned} \quad (22)$$

where the last inequality in (22) follows from the fact that $F_s^k - F_0^k \leq -\gamma c\varepsilon_f (\Delta^k)^p$ for successful iterations. Moreover, as in Case 1, $\Delta^{k+1} = \min\{\tau^{-1}\Delta^k, \delta_{\max}\}$ since the iteration is successful. Thus,

$$\mathbb{1}_{\bar{J}_k} \mathbb{1}_S (1 - \nu) [(\Delta^{k+1})^p - (\Delta^k)^p] \leq \mathbb{1}_{\bar{J}_k} \mathbb{1}_S (1 - \nu)(\tau^{-p} - 1)(\Delta^k)^p. \quad (23)$$

Then, choosing ν according to (11) yields

$$-\nu\gamma(\Delta^k)^p + (1 - \nu)(\tau^{-p} - 1)(\Delta^k)^p \leq 0. \quad (24)$$

Thus, combining (22), (23) and (24) leads to

$$\mathbb{1}_{\bar{J}_k} \mathbb{1}_S (\Phi_{k+1} - \Phi_k) \leq \mathbb{1}_{\bar{J}_k} \mathbb{1}_S \frac{\nu}{c\varepsilon_f} (|f(X^{k+1}) - F_s^k| + |F_0^k - f(X^k)|) \quad (25)$$

(ii) *Unsuccessful iteration* ($\mathbb{1}_{\bar{S}} = 1$). Δ^k is decreased and the change in function values is zero. Thus, the bound in the change of Φ_k follows straightforwardly from (18) by replacing $\mathbb{1}_{J_k}$ by $\mathbb{1}_{\bar{J}_k}$. More precisely, the following holds,

$$\begin{aligned} \mathbb{1}_{\bar{J}_k} \mathbb{1}_{\bar{S}} (\Phi_{k+1} - \Phi_k) &= -\mathbb{1}_{\bar{J}_k} \mathbb{1}_{\bar{S}} (1 - \nu)(1 - \tau^p) (\Delta^k)^p. \\ &\leq \mathbb{1}_{\bar{J}_k} \mathbb{1}_{\bar{S}} \frac{\nu}{c_{\mathcal{E}f}} (|f(X^{k+1}) - F_s^k| + |F_0^k - f(X^k)|) \end{aligned} \quad (26)$$

Then, combining (25) and (26), yields

$$\mathbb{1}_{\bar{J}_k} (\Phi_{k+1} - \Phi_k) \leq \mathbb{1}_{\bar{J}_k} \frac{\nu}{c_{\mathcal{E}f}} (|f(X^{k+1}) - F_s^k| + |F_0^k - f(X^k)|). \quad (27)$$

Taking conditional expectations with respect to \mathcal{F}_{k-1}^F in both sides of (27) and applying Lemma 1 leads to (21).

Now, combining expectations (13) and (21) leads to

$$\begin{aligned} \mathbb{E} (\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1}^F) &= \mathbb{E} ((\mathbb{1}_{J_k} + \mathbb{1}_{\bar{J}_k})(\Phi_{k+1} - \Phi_k) \mid \mathcal{F}_{k-1}^F) \\ &\leq [-\beta(1 - \nu)(1 - \tau^p) + 2\nu(1 - \beta)] (\Delta^k)^p. \end{aligned} \quad (28)$$

Then, choosing β according to (11) ensures that

$$-\beta(1 - \nu)(1 - \tau^p) + 2\nu(1 - \beta) \leq -\frac{1}{2}\beta(1 - \nu)(1 - \tau^p). \quad (29)$$

Hence, (12) follows from (28) and (29), which achieves the proof. \square

Summing both sides of (12) over $k \in \mathbb{N}$ and taking expectations with respect to \mathcal{F}_{k-1}^F lead to the following result generalizing that in [4], which shows in particular that the sequence $\{\Delta^k\}_{k \in \mathbb{N}}$ of step size parameters converges to zero with probability one.

Theorem 3. *Let all assumptions that were made in Theorem 2 hold. Then, the sequence $\{\Delta^k\}_{k \in \mathbb{N}}$ of step size parameters generated by Algorithm 1 satisfies for $p > 1$,*

$$\sum_{k=0}^{+\infty} (\Delta^k)^p < +\infty \quad \text{almost surely.}$$

Consider the stochastic process $\{(\Phi_k, \Delta^k, W_k)\}_{k \in \mathbb{N}}$, where Φ_k is the same random function in Theorem 2, Δ^k is the random step size parameter and $W_k = 2(\mathbb{1}_{J_k} - \frac{1}{2})$. Define $\hat{p} = \min(p - 1, 1)$ for some fixed $p > 1$. For some arbitrary fixed $\epsilon' \in (0, 1)$, consider the following random time $T_{\epsilon'}$ defined by

$$T_{\epsilon'} = \inf \left\{ k \in \mathbb{N} : \|\nabla f(X^k)\|^{1/\hat{p}} \leq \epsilon' \right\} \quad (30)$$

Then, $T_{\epsilon'}$ is a stopping time for the stochastic process generated by Algorithm 1 and is consequently a stopping time for $\{(\Phi_k, \Delta^k, W_k)\}_{k \in \mathbb{N}}$ [8, 14]. Moreover, $T_{\epsilon^{1/\hat{p}}}$ is the number of iterations required

by Algorithm 1 to drive the norm of the gradient of f below $\epsilon \in (0, 1)$. This latter remark will help to derive the main result of the present work in Theorem 4.

In order to apply Theorem 1 to $T_{\epsilon'}$, the remainder of this section is devoted to showing that Assumption 3 holds for the previous stochastic process. First, notice that since Theorem 2 holds without using any information about the existence of the gradient of f , then by multiplying both sides of (12) by $\mathbb{1}_{\{T_{\epsilon'} > k\}}$, Assumption 3-(iii), trivially holds with $\eta = \frac{1}{2}\beta(1 - \nu)(1 - \tau^p)$ and $h(x) = x^p$. By choosing λ such that $e^\lambda = \tau^{-1}$ and noticing that $\Delta^k \leq \delta_{\max} = \delta^0 e^{\lambda j_{\max}}$ in Algorithm 1 for all $k \in \mathbb{N}$, then Assumption 3-(i) holds.

Then, before showing that Assumption 3-(ii) also holds, let emphasize that as in the deterministic framework, polling directions in Algorithm 1 are chosen in such a way that their significant deterioration can be avoided asymptotically, i.e., in such a way to ensure that they never become close to losing the positive spanning property [15]. For this purpose, let recall the following definition of the *cosine measure* [10, 12] of a positive spanning set \mathbb{D}^k with non-zero vectors

$$\kappa(\mathbb{D}^k) := \min_{v \in \mathbb{R}^n} \max_{d \in \mathbb{D}^k} \frac{v^\top d}{\|v\| \|d\|}$$

In order to avoid the aforementioned deterioration of polling directions, the positive spanning sets are required to satisfy the following assumption [12, 15] where the size of the directions does not tend to infinite or approach zero, and the cosine measure always stays positive.

Assumption 5. *The followings hold for all positive spanning sets \mathbb{D}^k used for polling in Algorithm 1. There exists a constant $\kappa_{\min} > 0$ such that $\kappa(\mathbb{D}^k) > \kappa_{\min}$ for all k . There exist constants $d_{\min} > 0$ and $d_{\max} > 0$ such that $d_{\min} \leq \|d\| \leq d_{\max}$ for all $d \in \mathbb{D}^k$.*

The following result from [12] will be useful for the remaining of the analysis, and specifically the proof of the key result in Lemma 2. It shows by means of the cosine measure $\kappa(\mathbb{D}^k)$, how far can be in the worst case, the steepest descent direction, from the vector in \mathbb{D}^k which makes the smallest angle with $v = -\nabla f(x^k)$. This means in term of descent that, there exists $d_*^k \in \mathbb{D}^k$ such that

$$\kappa(\mathbb{D}^k) \|\nabla f(x^k)\| \|d_*^k\| \leq -\nabla f(x^k)^\top d_*^k. \quad (31)$$

For the remaining of the analysis, the following standard assumption is also imposed on the gradient of f .

Assumption 6. *The gradient ∇f of the objective function f is L -Lipschitz continuous everywhere.*

Then, define the constant $\delta_{\epsilon'}$ as follows

$$\delta_{\epsilon'} = \frac{\epsilon'}{\zeta} \quad \text{with} \quad \zeta > \left[\kappa_{\min}^{-1} (L d_{\max} + (\gamma + 2) c \varepsilon_f d_{\min}^{-1}) \right]^{1/\hat{p}}, \quad (32)$$

where without loss of generality, $L d_{\max} > \kappa_{\min}$ so that $\delta_{\epsilon'} < 1$ for the needs of the analysis and specifically, the proof of Lemma 2. Then following [8], it can be assumed without any loss of generality that $\delta_{\epsilon'} = \tau^{-i} \delta^0$, for some integer $i \leq 0$. Hence, for any k , $\Delta^k = \tau^{i_k} \delta_{\epsilon'}$, for some integer i_k . Thus, what remains to be proved in Assumption 3 in order to apply Theorem 1 is the dynamics (8). Note

however that the proof of the dynamics (8) which will be achieved in Lemma 3, need the following intermediate key result. Indeed, in the stochastic trust region framework of [8], the proof of a similar dynamics strongly relies on the fact that any iteration k , where $\|\nabla f(x_k)\| > \epsilon$ and for which “good” model and estimates occur, is successful provided that the *trust region radius* δ_k is below a threshold Δ_ϵ . Nevertheless, unlike the trust region framework where informations can possibly easily be derived on the true gradient $\nabla f(x^k)$ using those provided by the *gradient estimate* g_k , the algorithmic framework of the present work does not use any gradient information. Thus, the main challenge in proving that Assumption 3-(ii) holds, lies in linking the event $\{\|\nabla f(X^k)\|^{1/\hat{p}} > \epsilon'\}$ to a successful iteration of Algorithm 1, which is done next.

Lemma 2. *Assume that Assumption 6 and 5 hold and that $\delta^k \leq \delta_{\epsilon'}$. Let f_0^k and f_s^k be ϵ_f -accurate estimates of $f(x^k)$ and $f(x^k + s^k)$ respectively. If $\|\nabla f(x^k)\|^{1/\hat{p}} > \epsilon'$, then*

$$f_s^k - f_0^k \leq -\gamma c \epsilon_f (\delta^k)^p.$$

In particular, this means that the iteration k of Algorithm 1 is successful.

Proof. The proof uses elements derived in [12]. Suppose that $\delta^k \leq \delta_{\epsilon'}$ and assume in contradiction that $f_s^k - f_0^k > -\gamma c \epsilon_f (\delta^k)^p$. Since the estimates f_0^k and f_s^k are ϵ_f -accurate, then it follows from the following equality

$$f(x^k + s^k) - f(x^k) = f(x^k + s^k) - f_s^k + (f_s^k - f_0^k) + f_0^k - f(x^k)$$

that

$$f(x^k + s^k) - f(x^k) + (\gamma + 2)c \epsilon_f (\delta^k)^p \geq 0. \quad (33)$$

Recall that $s^k = \delta^k d$ where $d \in \mathbb{D}^k$ denotes any direction used by Algorithm 1 at iteration k . It follows from the mean value theorem, combined with (33), that there exists a constant $\mu_k \in [0, 1]$ such that

$$0 \leq \delta^k \nabla f(x^k + \mu_k \delta^k d_*^k)^\top d_*^k + (\gamma + 2)c \epsilon_f (\delta^k)^p, \quad (34)$$

where d_*^k is the direction satisfying (31). Dividing both sides of (34) by δ^k and subtracting $\nabla f(x^k)^\top d_*^k$, yields

$$-\nabla f(x^k)^\top d_*^k \leq [\nabla f(x^k + \mu_k \delta^k d_*^k) - \nabla f(x^k)]^\top d_*^k + (\gamma + 2)c \epsilon_f (\delta^k)^{p-1}. \quad (35)$$

Putting (31) and (35) together, yields

$$\kappa(\mathbb{D}^k) \|\nabla f(x^k)\| \|d_*^k\| \leq [\nabla f(x^k + \mu_k \delta^k d_*^k) - \nabla f(x^k)]^\top d_*^k + (\gamma + 2)c \epsilon_f (\delta^k)^{p-1}. \quad (36)$$

Then, dividing both sides of (36) by $\kappa(\mathbb{D}^k) \|d_*^k\|$ and using Assumption 6 and 5, lead to

$$\begin{aligned} \|\nabla f(x^k)\| &\leq \kappa_{\min}^{-1} [L d_{\max} \delta^k + (\gamma + 2)c \epsilon_f d_{\min}^{-1} (\delta^k)^{p-1}] \\ &\leq \kappa_{\min}^{-1} (L d_{\max} + (\gamma + 2)c \epsilon_f d_{\min}^{-1}) (\delta^k)^{\min(p-1, 1)}, \end{aligned} \quad (37)$$

where the inequality (37) follows from the fact that $\delta^k \leq \delta_{\epsilon'} < 1$. Now, recall that $\hat{p} = \min(p-1, 1)$ and let $L_1 := \kappa_{\min}^{-1} (L d_{\max} + (\gamma + 2)c \epsilon_f d_{\min}^{-1})$. Then, it follows from (37) that

$$\|\nabla f(x^k)\|^{1/\hat{p}} \leq L_1^{1/\hat{p}} \delta^k \leq L_1^{1/\hat{p}} \delta_{\epsilon'} = L_1^{1/\hat{p}} \frac{\epsilon'}{\zeta} \leq \epsilon', \quad (38)$$

where the last inequality in (38) follows from (32), which achieves the proof. \square

Finally, the following result shows that the dynamics (8) of Assumption 3-(ii) holds.

Lemma 3. *Let Assumption 6 and all assumptions that were made in Theorem 2 hold. Then Assumption 3-(ii) is satisfied for the random variable $W_k = 2(\mathbb{1}_{J_k} - \frac{1}{2})$, $\lambda = -\ln(\tau)$ and $q = \beta$.*

Proof. The result is proved by adapting the proof of a similar Lemma from [8]. First, notice that (8) trivially holds when $\mathbb{1}_{\{T_{\epsilon'} > k\}} = 0$. Thus, the remaining of the proof is devoted to showing that conditioned on the event $\{T_{\epsilon'} > k\}$, i.e., when $\mathbb{1}_{\{T_{\epsilon'} > k\}} = 1$, then the following holds

$$\Delta^{k+1} \geq \min \{ \delta_{\epsilon'}, \min \{ \tau^{-1} \Delta^k, \delta_{\max} \} \mathbb{1}_{J_k} + \tau \Delta^k \mathbb{1}_{\bar{J}_k} \}. \quad (39)$$

Notice that every realization such that $\delta^k > \delta_{\epsilon'}$ also satisfies $\delta^k \geq \tau^{-1} \delta_{\epsilon'}$ whence $\delta^{k+1} \geq \tau \delta^k \geq \delta_{\epsilon'}$. Now, assume that $\delta^k \leq \delta_{\epsilon'}$. Since $T_{\epsilon'} > k$, then it is the case that $\|\nabla f(x^k)\|^{1/\hat{p}} > \epsilon'$. If $\mathbb{1}_{J_k} = 1$, then the estimates are good and are specifically ε_f -accurate. Hence, it follows from Lemma 2 that the k th iteration is successful. Thus, $x^{k+1} = x^k + s^k$ and $\delta^{k+1} = \min \{ \tau^{-1} \delta^k, \delta_{\max} \}$. But if $\mathbb{1}_{J_k} = 0$, i.e., $\mathbb{1}_{\bar{J}_k} = 1$, then the inequality $\delta^{k+1} \geq \tau \delta^k$ always holds by the dynamics of Algorithm 1. The proof is complete by noticing finally that $\mathbb{P}(J_k | \mathcal{F}_{k-1}^F) \geq q = \beta$. \square

4.2 Complexity result and first-order optimality conditions

The following result provides a bound on the expected number of iterations taken by Algorithm 1 before $\{\|\nabla f(X^k)\| \leq \epsilon\}$ occurs and is the main result of the present work.

Theorem 4. *Let Assumption 6 and all assumptions that were made in Theorem 2 hold with $\beta \in (1/2, 1)$ and $\nu \in (0, 1)$ satisfying (11). Consider Algorithm 1 and the corresponding stochastic process. For some arbitrary fixed $\epsilon \in (0, 1)$, consider the random time T_ϵ^* defined by*

$$T_\epsilon^* = \inf \{ k \in \mathbb{N} : \|\nabla f(X^k)\| \leq \epsilon \}. \quad (40)$$

Then,

$$\mathbb{E}(T_\epsilon^*) \leq \frac{2\Phi_0 L_2}{(2\beta - 1)(1 - \nu)(1 - \tau^p)} \epsilon^{\frac{-p}{\min(p-1, 1)}} + 1, \quad (41)$$

where $L_2 := [1 + \kappa_{\min}^{-1} (Ld_{\max} + (\gamma + 2)c\varepsilon_f d_{\min}^{-1})]^{\frac{p}{\min(p-1, 1)}}$. i.e., the expected number of iterations taken by Algorithm 1 to reduce the gradient below $\epsilon \in (0, 1)$ is bounded in $\mathcal{O}\left(\epsilon^{\frac{-p}{\min(p-1, 1)}} / (2\beta - 1)\right)$.

Proof. As shown previously, since Assumption 3 holds for the stochastic process $\{(\Phi_k, \Delta^k, W_k)\}_{k \in \mathbb{N}}$ generated by Algorithm 1, with $q = \beta$, $h(x) = x^p$, $\eta = \frac{1}{2}\beta(1 - \nu)(1 - \tau^p)$ and $\delta_{\epsilon'} = \epsilon'/\zeta$, then Theorem 1 applies for the stopping time $T_{\epsilon'}$ defined in (30). Thus, the following inequality holds for all $\epsilon' \in (0, 1)$

$$\mathbb{E}(T_{\epsilon'}) \leq \frac{\beta}{2\beta - 1} \times \frac{\Phi_0 \zeta^p}{\eta \epsilon'^p} + 1, \quad (42)$$

where $\zeta^p > [\kappa_{\min}^{-1} (Ld_{\max} + (\gamma + 2)c\varepsilon_f d_{\min}^{-1})]^{p/\hat{p}}$ thanks to (32), with $\hat{p} = \min(p - 1, 1)$. Now, let $\epsilon \in (0, 1)$ be arbitrary fixed. Then, $\epsilon^{1/\hat{p}} \in (0, 1)$, which means that (42) holds in particular for $\epsilon' = \epsilon^{1/\hat{p}}$. By noticing moreover that $T_{\epsilon^{1/\hat{p}}} = T_\epsilon^*$, then it follows from (42) that

$$\mathbb{E}(T_\epsilon^*) \leq \frac{\beta}{2\beta - 1} \times \frac{2\Phi_0 \zeta^p}{\beta(1 - \nu)(1 - \tau^p)} \epsilon^{-p/\hat{p}} + 1. \quad (43)$$

Since $p/\hat{p} = p/\min(p-1, 1) \geq 2$ for all $p > 1$, then (41) results from (43) by choosing ζ^p according to $\zeta^p = [1 + \kappa_{\min}^{-1} (Ld_{\max} + (\gamma + 2)c\varepsilon_f d_{\min}^{-1})]^{p/\hat{p}}$, which achieves the proof. \square

The following lim inf-type first-order necessary optimality condition is a simple consequence of the complexity result of Theorem 4. It shows the existence of a subsequence of random iterates generated by Algorithm 1 which drives the norm of the gradient of f to zero with probability one. Note that a similar corollary has been derived in [14].

Theorem 5. *Let Assumption 6 and all assumptions that were made in Theorem 2 hold. Then the sequence $\{X^k\}_{k \in \mathbb{N}}$ of random iterates generated by Algorithm 1 satisfies*

$$\liminf_{k \rightarrow +\infty} \|\nabla f(X^k)\| = 0 \quad \text{almost surely.} \quad (44)$$

Discussion

This manuscript presents the first convergence rate analysis of a broad class of stochastic directional direct-search (SDDS) algorithms, designed for the unconstrained optimization of noisy blackboxes, and based on imposing a sufficient decrease condition when accepting new iterates. Using an existing supermartingale-based framework for the analysis, the methodology for deriving the worst case complexity of SDDS algorithms heavily relies on bounding an expected stopping time associated to the stochastic process generated by the algorithms. The analysis showed that SDDS algorithms have the same worst case complexity as any other first-order optimization method in a nonconvex setting. In particular, this complexity bound matches in some sense its deterministic counterparts despite the fact that function estimates are sometimes allowed to be arbitrarily inaccurate. The main novelty of the present research compared to many others on the worst case complexity analysis of stochastic DFO methods, lies in the fact that the proposed method does not need any gradient information to find descent directions.

The analysis in the present manuscript strongly relies on the assumption that function estimates are unbiased. Thus, obtaining worst case complexity results when such estimates are possibly biased is a topic for future research.

Acknowledgments

The author is grateful to Sébastien Le Digabel and Charles Audet from Polytechnique Montréal and Michael Kokkolaras from McGill university for valuable discussions and constructive suggestions that improve the quality of the presentation. This work is supported by the NSERC CRD RDCPJ 490744-15 grant and by an InnovÉÉ grant, both in collaboration with Hydro-Québec and Rio Tinto.

References

- [1] S. Alarie, C. Audet, P. Y. Bouchet, and S. Le Digabel. Optimization of noisy blackboxes with adaptive precision. Technical Report G-2019-84, Les cahiers du GERAD, 2019.

- [2] C. Audet and J.E. Dennis, Jr. Analysis of generalized pattern searches. *SIAM Journal on Optimization*, 13(3):889–903, 2003.
- [3] C. Audet and J.E. Dennis, Jr. Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006.
- [4] C. Audet, K. J. Dzahini, M. Kokkolaras, and S. Le Digabel. StoMADS: Stochastic blackbox optimization using probabilistic estimates. Technical Report G-2019-30, Les cahiers du GERAD, 2019.
- [5] C. Audet and W. Hare. *Derivative-Free and Blackbox Optimization*. Springer Series in Operations Research and Financial Engineering. Springer International Publishing, Cham, Switzerland, 2017.
- [6] A. S. Berahas, L. Cao, and K. Scheinberg. Global Convergence Rate Analysis of a Generic Line Search Algorithm with Noise. *arXiv*, 2019.
- [7] R.N. Bhattacharya and E.C. Waymire. *A basic course in probability theory*, volume 69. Springer, 2007.
- [8] J. Blanchet, C. Cartis, M. Menickelly, and K. Scheinberg. Convergence Rate Analysis of a Stochastic Trust-Region Method via Supermartingales. *INFORMS Journal on Optimization*, 2019.
- [9] R. Chen, M. Menickelly, and K. Scheinberg. Stochastic optimization using a trust-region method and random models. *Mathematical Programming*, 169(2):447–487, 2018.
- [10] A.R. Conn, K. Scheinberg, and L.N. Vicente. *Introduction to Derivative-Free Optimization*. MOS-SIAM Series on Optimization. SIAM, Philadelphia, 2009.
- [11] S. Gratton, C.W. Royer, L.N. Vicente, and Z. Zhang. Direct search based on probabilistic descent. *SIAM Journal on Optimization*, 25(3):1515–1541, 2015.
- [12] T.G. Kolda, R.M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45(3):385–482, 2003.
- [13] J. Larson and S.C. Billups. Stochastic derivative-free optimization using a trust region framework. *Computational Optimization and Applications*, 64(3):619–645, 2016.
- [14] C. Paquette and K. Scheinberg. A Stochastic Line Search Method with Expected Complexity Analysis. *SIAM Journal on Optimization*, 30(1):349–376, 2020.
- [15] L. N. Vicente. Worst case complexity of direct search. *EURO Journal on Computational Optimization*, 1(1):143–153, 2013.
- [16] X. Wang and Y. Yuan. Stochastic Trust Region Methods with Trust Region Radius Depending on Probabilistic Models. *arXiv*, 2019.