



Noisy Optimization: Convergence with a Fixed Number of Resamplings

Marie-Liesse Cauwet

► To cite this version:

Marie-Liesse Cauwet. Noisy Optimization: Convergence with a Fixed Number of Resamplings. EvoStar, Apr 2014, Granada, Spain. hal-00976063

HAL Id: hal-00976063

<https://inria.hal.science/hal-00976063>

Submitted on 9 Apr 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Noisy Optimization: Convergence with a Fixed Number of Resamplings

Marie-Liesse Cauwet

TAO (Inria), LRI, UMR 8623 (CNRS - Univ. Paris-Sud), France

Abstract. It is known that evolution strategies in continuous domains might not converge in the presence of noise [3, 14]. It is also known that, under mild assumptions, and using an increasing number of resamplings, one can mitigate the effect of additive noise [4] and recover convergence. We show new sufficient conditions for the convergence of an evolutionary algorithm with constant number of resamplings; in particular, we get fast rates (log-linear convergence) provided that the variance decreases around the optimum slightly faster than in the so-called multiplicative noise model.

Keywords: Noisy optimization, evolutionary algorithm, theory.

1 Introduction

Given a domain $\mathcal{D} \in \mathbb{R}^d$, with d a positive integer, a noisy objective function is a stochastic process $f : (x, \omega) \mapsto f(x, \omega)$ with $x \in \mathcal{D}$ and ω a random variable independently sampled at each call to f . Noisy optimization is the search of x such that $\mathbb{E}[f(x, \omega)]$ is approximately minimum. Throughout the paper, x^* denotes the unknown exact optimum, supposed to be unique. For any positive integer n , \tilde{x}_n denotes the search point used in the n^{th} function evaluation. We here consider black-box noisy optimization, i.e we can have access to f only through calls to a black-box which, on request x , (i) randomly samples ω (ii) returns $f(x, \omega)$. Among zero-order methods proposed to solve noisy optimization problems, some of the most usual are evolution strategies; [1] has studied the performance of evolution strategies in the presence of noise, and investigated its robustness by tuning the population size of the offspring and the mutation strength. Another approach consists in using resamplings of each individual (averaging multiple resamplings reduces the noise), rather than increasing the population size. Resampling means that, when evaluating $f(x, \omega)$, several independent copies $\omega_1, \dots, \omega_r$ of ω are used (i.e. the black-box oracle is called several times with a same x) and we use as an approximate fitness value $\frac{1}{r} \sum_{i=1}^r f(x, \omega_i)$ in the optimization algorithm. The key point is how to choose r , number of resamplings, for a given x . Another crucial point is the model of noise. Different models of noise can be considered: additive noise (Eq. 3), multiplicative noise (Eq. 4) or a more general model (Eq. 5). Notice that, in Eq. 5 when $z > 0$, the noise decreases to zero near the optimum; this setting is not artificial as we can observe this behavior in many real problems.

Let us give an example in which the noise variance decreases to zero around the optimum. Consider a Direct Policy Search problem, i.e. the optimization of a parametric policy on simulations. Assume that we optimize the success rate of a policy. Assume that the optimum policy has a success rate 100%. Then, the variance is zero at the optimum.

1.1 Convergence Rates: log-linear convergence and log-log convergence

Depending on the specific class of optimization problems and on some internal properties of the algorithm considered, we obtain different uniform rates of convergence (where the convergence can be almost sure, in probability or in expectation, depending on the setting); a fast rate will be a log-linear convergence, as follows:

$$\textbf{Fast rate: } \limsup_n \frac{\log \|\tilde{x}_n - x^*\|}{n} = -A < 0, \quad (1)$$

In the noise-free case, evolution strategies typically converge linearly in log-linear scale, as shown in [5, 7, 8, 15, 18].

The algorithm presents a slower rate of convergence in case of log-log convergence, as follows:

$$\textbf{Slow rate: } \limsup_n \frac{\log \|\tilde{x}_n - x^*\|}{\log n} = -A < 0, \quad (2)$$

The log-log rates are typical rates in the noisy case (see [2, 4, 9–11, 16, 17]). Nevertheless, we will here show that, under specific assumptions on the noise (if the noise around the optimum decreases “quickly enough”, see section 1.4), we can reach faster rates: log-linear convergence rates as in Eq. 1, by averaging a constant number of resamplings of $f(x, \omega)$.

1.2 Additive noise model

Additive noise refers to:

$$f(x, \omega) = \|x - x^*\|^p + \textit{noise}_\omega, \quad (3)$$

where p is a positive integer and where \textit{noise}_ω is sampled independently with a fixed given distribution. In this model, the noise has lower bounded variance, even in the neighborhood of the optimum. The uniform rate typically converges linearly in log – log scale (cf Eq. 2) as discussed in [2, 9–11, 16, 17]. This important case in applications has been studied in [9, 11, 12, 16] where tight bounds have been shown for stochastic gradient algorithms using finite differences. When using evolution strategies, [4] has shown mathematically that an exponential number of resamplings (number of resamplings scaling exponentially with the index of iterations) or an adaptive number of resamplings (scaling as a polynomial of the inverse step-size) can both lead to a log-log convergence rate.

1.3 Multiplicative noise model

Multiplicative noise, in the unimodal spherical case, refers to

$$f(x, \omega) = \|x - x^*\|^p + \|x - x^*\|^p \times \text{noise}_\omega \quad (4)$$

and some compositions (by increasing mappings) of this function, where p is a positive integer and where noise_ω is sampled independently with a fixed given distribution. [14] has studied the convergence of evolution strategies in noisy environments with multiplicative noise, and essentially shows that the result depends on the noise distribution: if noise_ω is conveniently lower bounded, then some standard $(1 + 1)$ evolution strategy converges to the optimum; if arbitrarily negative values can be sampled with non-zero probability, then it does not converge.

1.4 A more general noise model

Eqs. 3 and 4 are particular cases of a more general noise model:

$$f(x, \omega) = \|x - x^*\|^p + \|x - x^*\|^{pz/2} \times \text{noise}_\omega. \quad (5)$$

where p is a positive integer, $z \geq 0$ and noise_ω is sampled independently with a fixed given distribution. Eq. 5 boils down to Eq. 3 when $z = 0$ and to Eq. 4 when $z = 2$. We will here obtain fast rates for some larger values of z . More precisely, we will show that when $z > 2$, we obtain log-linear rates, as in Eq. 1. Incidentally, this shows some tightness (with respect to z) of conditions for non-convergence in [14].

2 Theoretical analysis

Section 2.1 is devoted to some preliminaries. Section 2.2 presents results for constant numbers of resamplings on our generalized noise model (Eq. 5) when $z > 2$.

2.1 Preliminary: noise-free case

Typically, an evolution strategy at iteration n :

- generates λ individuals using the current estimate x_{n-1} of the optimum x^* and the so-called mutation strength (or step-size) σ_{n-1} ,
- provides a pair (x_n, σ_n) where x_n is a new estimate of x^* and σ_n is a new mutation strength.

From now on, for the sake of notation simplicity, we assume that $x^* = 0$.

For some evolution strategies and in the noise-free case, we know (see e.g. Theorem 4 in [5]) that there exists a constant A such that :

$$\frac{\log(\|x_n\|)}{n} \xrightarrow[n \rightarrow \infty]{a.s.} -A \quad (6)$$

$$\frac{\log(\sigma_n)}{n} \xrightarrow[n \rightarrow \infty]{a.s.} -A \quad (7)$$

This paper will discuss cases in which an algorithm verifying Eqs. 6, 7 in the noise-free case also verifies them in a noisy setting.

Remarks: *In the general case of arbitrary evolution strategies (ES), we don't know if A is positive, but:*

- *in the case of a $(1+1)$ -ES with generalized one-fifth success rule, $A > 0$ see [6];*
- *in the case of a self-adaptive $(1, \lambda)$ -ES with gaussian mutations, the estimate of A by Monte-Carlo simulations is positive [5].*

Property 1. For some $\delta > 0$, for any α, α' such that $\alpha < A$ and $\alpha' > A$, there exist $C > 0, C' > 0, V > 0, V' > 0$, such that with probability at least $1 - \delta$

$$\forall n \geq 1, C' \exp(-\alpha' n) \leq \|x_n\| \leq C \exp(-\alpha n); \quad (8)$$

$$\forall n \geq 1, V' \exp(-\alpha' n) \leq \sigma_n \leq V \exp(-\alpha n). \quad (9)$$

Proof. For any $\alpha < A$, almost surely, $\log(\|x_n\|) \leq -\alpha n$ for n sufficiently large. So, almost surely, $\sup_{n \geq 1} \log(\|x_n\|) + \alpha n$ is finite. Consider V the quantile $1 - \frac{\delta}{4}$ of $\exp(\sup_{n \geq 1} \log(\|x_n\|) + \alpha n)$. Then, with probability at least $1 - \frac{\delta}{4}$, $\forall n \geq 1, \|x_n\| \leq V \exp(-\alpha n)$. We can apply the same trick for lower bounding $\|x_n\|$, and upper and lower bounding σ_n , all of them with probability $1 - \frac{\delta}{4}$, so that all bounds hold true simultaneously with probability at least $1 - \delta$. \square

2.2 Noisy case

The purpose of this Section is to show that if some evolution strategies perform well (linear convergence in the log-linear scale, as in Eqs. 6, 7), then, just by considering Y resamplings for each fitness evaluation as explained in Alg. 1, they will also be fast in the noisy case.

Our theorem holds for any evolution strategy satisfying the following constraints:

- At each iteration n , a search point x_n is defined and λ search points are generated and have their fitness values evaluated.
- The noisy fitness values are averaged over Y (a constant) resamplings.
- The j^{th} individual evaluated at iteration n is randomly drawn by $x_n + \sigma_n \mathcal{N}_d$ with \mathcal{N}_d a d -dimensional standard Gaussian variable.

This framework is presented in Alg. 1.

We now state our theorem, under log-linear convergence assumption (cf assumption (ii) below).

Algorithm 1 A general framework for evolution strategies. For simplicity, it does not cover all evolution strategies, e.g. mutations of step-sizes as in self-adaptive algorithms are not covered; yet, our proof can be extended to a more general case ($x_{n,i}$ distributed as $x_n + \sigma N$ for some noise N with exponentially decreasing tail). The case $Y = 1$ is the case without resampling. Our theorem basically shows that if such an algorithm converges linearly (in log-linear scale) in the noise-free case then the version with Y large enough converges linearly in the noisy case when $z > 2$.

```

Initialize  $x_0$  and  $\sigma_0$ .
 $n \leftarrow 1$ 
while not finished do
  for  $i \in \{1, \dots, \lambda\}$  do
    Define  $x_{n,i} = x_n + \sigma_n \mathcal{N}_d$ .
    Define  $y_{n,i} = \frac{1}{Y} \sum_{k=1}^Y f(x_{n,i}, \omega_k)$ .
  end for
  Update:  $(x_{n+1}, \sigma_{n+1}) \leftarrow \text{update}(x_{n,1}, \dots, x_{n,\lambda}, y_{n,1}, \dots, y_{n,\lambda}, \sigma_n)$ .
   $n \leftarrow n + 1$ 
end while

```

Theorem 1. Consider the following assumptions:

(i) the fitness function f satisfies $\mathbb{E}[f(x, \omega)] = \|x\|^p$ and has a limited variance:

$$\text{Var}(f(x, \omega)) \leq (\mathbb{E}[f(x, \omega)])^z \text{ for some } z > 2; \quad (10)$$

(ii) in the noise-free case, the ES with population size λ under consideration is log-linearly converging, i.e. for any $\delta > 0$, for some $\alpha > 0$, $\alpha' > 0$, there exist $C > 0$, $C' > 0$, $V > 0$, $V' > 0$, such that with probability $1 - \delta$, Eqs. 8 and 9 hold;

(iii) the number Y of resamplings per individual is constant.

Then, if $z > \max\left(\frac{2(p\alpha' - (\alpha - \alpha')d)}{p\alpha}, \frac{2(2\alpha' - \alpha)}{\alpha}\right)$, for any $\delta > 0$, there is $Y_0 > 0$ such that for any $Y \geq Y_0$, Eqs. 8 and 9 also hold with probability at least $(1 - \delta)^2$ in the noisy case.

Corollary 1. Under the same assumptions, with probability at least $(1 - \delta)^2$,

$$\limsup_n \frac{\log(\|\tilde{x}_n\|)}{n} \leq -\frac{\alpha}{\lambda Y}$$

Proof of Corollary 1 : Immediate consequence of Theorem 1, by applying Eq. 8 and using $\limsup_n \frac{\log(\|\tilde{x}_n\|)}{n} = \limsup_n \frac{\log(\|x_n\|)}{\lambda Y n}$. \square

Remarks:

- **Interpretation:** Informally speaking, our theorem shows that if an algorithm converges in the noise-free case, then it also converges in the noisy case with the resampling rule, at least if z and Y are large enough.

- Notice that we can choose constants α and α' very close to each other. Then the assumption $z > \max\left(\frac{2(p\alpha' - (\alpha - \alpha')d)}{p\alpha}, \frac{2(2\alpha' - \alpha)}{\alpha}\right)$ boils down to $z > 2$.
- We show a log-linear convergence rate as in the noise-free case. This means that we get $\log \|\tilde{x}_n\|$ linear in the number of function evaluations. This is as Eq. 1, and faster than Eq. 2 which is typical for noisy optimization with constant variance.
- In the previous hypothesis, the new individuals are drawn following $x_n + \sigma_n \mathcal{N}_d$ with \mathcal{N}_d a d -dimensional standard Gaussian variable, but we could substitute \mathcal{N}_d for any random variable with an exponentially decreasing tail.

Proof of Theorem 1 : In all the proof, \mathcal{N}_k denotes a standard normal random variable in dimension k .

Sketch of proof: Consider an arbitrary $\delta > 0$ and $\delta_n = \exp(-\gamma n)$ for some $n \geq 1$ and $\gamma > 0$.

We compute in Lemma 2 the probability that at least two generated points x_{n,i_1} and x_{n,i_2} at iteration n are “close”, i.e are such that $||x_{n,i_1}|^p - |x_{n,i_2}|^p| \leq \delta_n$; then we calculate the probability that the noise of at least one of the λ evaluated individuals of iteration n is bigger than $\frac{\delta_n}{2}$ in Lemma 3. Thus, we can conclude in Lemma 4 by estimating the probability that at least two individuals are misranked due to noise.

We first begin by showing a technical lemma.

Lemma 1. *Let $u \in \mathbb{R}^d$ be a unit vector and \mathcal{N}_d a d -dimensional standard normal random variable. Then for $S > 0$ and $\ell > 0$, there exists a constant $M > 0$ such that :*

$$\max_{v \geq 0} \mathbb{P}(| |u + S\mathcal{N}_d|^p - v | \leq \ell) \leq MS^{-d} \max\left(\ell, \ell^{d/p}\right).$$

Proof. For any $v \geq \ell$, we denote $E_{v \geq \ell}$ the set :

$$E_{v \geq \ell} = \{x ; | |x|^p - v | \leq \ell\} = \left\{x ; (v - \ell)^{\frac{1}{p}} \leq |x| \leq (v + \ell)^{\frac{1}{p}}\right\}.$$

We first compute $\mu(E_{v \geq \ell})$, the Lebesgue measure of $E_{v \geq \ell}$:

$$\mu(E_{v \geq \ell}) = K_d \left\{ (v + \ell)^{\frac{d}{p}} - (v - \ell)^{\frac{d}{p}} \right\},$$

with $K_d = \frac{(2\pi)^{d/2}}{2 \times 4 \times \dots \times d}$ if d is even, and $K_d = \frac{2(2\pi)^{(d-1)/2}}{1 \times 3 \times \dots \times d}$ otherwise. Hence, by Taylor expansion, $\mu(E_{v \geq \ell}) \leq K v^{\frac{d}{p}-1} \ell$, where $K = K_d \left(2 \frac{d}{p} + \sup_{v \geq \ell} \sup_{0 < \zeta < \frac{\ell}{v}} \frac{q''(\zeta)}{2} \frac{\ell}{v} \right)$,

with $q(x) = (1 + x)^{\frac{d}{p}}$.

• If $v \geq \ell$:

$$\begin{aligned} \mathbb{P}(| |u + S\mathcal{N}_d|^p - v | \leq \ell) &= \mathbb{P}(u + S\mathcal{N}_d \in E_{v \geq \ell}), \\ &\leq S^{-d} \sup_{x \in E_{v \geq \ell}} \left(\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\|S^{-1}(x - u)\|^2}{2}\right) \right) \mu(E_{v \geq \ell}), \\ &\leq M_1 S^{-d} \ell, \\ &\leq M_1 S^{-d} \max\left(\ell, \ell^{d/p}\right). \end{aligned}$$

where $M_1 = \frac{K}{\sqrt{2\pi}} \sup_{v \geq \ell} \sup_{x: \|x\| \leq (v+\ell)^{\frac{1}{p}}} \left[v^{\frac{d}{p}-1} \exp\left(-\frac{\|S^{-1}(x-u)\|^2}{2}\right) \right]$.

- If $v < \ell$, $\mathbb{P}(|\|u + S\mathcal{N}_d\|^p - v| \leq \ell) \leq M_2 S^{-d} \ell^{d/p} \leq M_2 S^{-d} \max(\ell, \ell^{d/p})$,

where $M_2 = 2^{\frac{d}{p}} \frac{K_d}{\sqrt{2\pi}}$. Hence the result follows by taking $M = \max(M_1, M_2)$. \square

Lemma 2. Let us denote by $P_n^{(1)}$ the probability that, at iteration n , there exist at least two points x_{n,i_1} and x_{n,i_2} such that $|\|x_{n,i_1}\|^p - \|x_{n,i_2}\|^p| \leq \delta_n$. Then

$$P_n^{(1)} \leq B\lambda^2 \exp(-\gamma'n),$$

for some $B > 0$ and $\gamma' > 0$ depending on $\gamma, d, p, C, C', V, \alpha, \alpha'$.

Proof. Let us first compute the probability $P_n^{(0)}$ that, at iteration n , two given generated points x_{n,i_1} and x_{n,i_2} are such that $|\|x_{n,i_1}\|^p - \|x_{n,i_2}\|^p| \leq \delta_n$. Let us denote by \mathcal{N}_d^1 and \mathcal{N}_d^2 two d -dimensional standard independent random variables, $u \in \mathbb{R}^d$ a unit vector and $S_n = \frac{\sigma_n}{\|x_n\|}$.

$$\begin{aligned} P_n^{(0)} &= \mathbb{P}(|\|x_n + \sigma_n \mathcal{N}_d^1\|^p - \|x_n + \sigma_n \mathcal{N}_d^2\|^p| \leq \delta_n), \\ &= \mathbb{P}\left(|\|u + S_n \mathcal{N}_d^1\|^p - \|u + S_n \mathcal{N}_d^2\|^p| \leq \frac{\delta_n}{\|x_n\|^p}\right), \\ &\leq \max_{v \geq 0} \mathbb{P}\left(|\|u + S_n \mathcal{N}_d^1\|^p - v| \leq \frac{\delta_n}{\|x_n\|^p}\right). \end{aligned}$$

Hence, by Lemma 1, there exists a $M > 0$ such that $P_n^{(0)} \leq M S_n^{-d} \left(\frac{\delta_n}{\|x_n\|^p}\right)^m$,

where m is such that $\left(\frac{\delta_n}{\|x_n\|^p}\right)^m = \max\left(\frac{\delta_n}{\|x_n\|^p}, \left(\frac{\delta_n}{\|x_n\|^p}\right)^{d/p}\right)$. Moreover $S_n \geq V' C^{-1} \exp(-(\alpha' - \alpha)n)$ by Assumption (ii). Thus $P_n^{(0)} \leq B \exp(-\gamma'n)$, with $B = M V'^{-d} C^d C'^{-mp}$ and $\gamma' = d(\alpha - \alpha') + m\gamma - mp\alpha'$. In particular, γ' is positive, provided that γ is sufficiently large.

By union bound, $P_n^{(1)} \leq \frac{(\lambda-1)\lambda}{2} P_n^{(0)} \leq B\lambda^2 \exp(-\gamma'n)$. \square

We now provide a bound on the probability $P_n^{(3)}$ that the fitness value of at least one search point generated at iteration n has noise (i.e. deviation from expected value) bigger than $\frac{\delta_n}{2}$ in spite of the Y resamplings.

Lemma 3.

$$\begin{aligned} P_n^{(3)} &:= \mathbb{P}\left(\exists i \in \{1, \dots, \lambda\} ; \left| \frac{1}{Y} \sum_{j=1}^Y f(x_{n,i}, \omega_j) - \mathbb{E}[f(x_{n,i}, \omega_j)] \right| \geq \frac{\delta_n}{2}\right) \\ &\leq \lambda B' \exp(-\gamma''n) \end{aligned}$$

for some $B' > 0$ and $\gamma'' > 0$ depending on $\gamma, d, p, z, C, Y, \alpha, \alpha'$.

Proof. First, for one point x_{n,i_0} , $i_0 \in \{1, \dots, \lambda\}$ generated at iteration n , we write $P_n^{(2)}$ the probability that when evaluating the fitness function at this point, we make a mistake bigger than $\frac{\delta_n}{2}$.

$P_n^{(2)} = \mathbb{P}(|\frac{1}{Y} \sum_{j=1}^Y f(x_{n,i_0}, \omega_j) - \mathbb{E}[f(x_{n,i_0}, \omega_j)]| \geq \frac{\delta_n}{2}) \leq B' \exp(-\gamma'' n)$ by using Chebyshev's inequality, where $B' = 4Y^{-1}C^{pz}$ and $\gamma'' = \alpha z p - 2\gamma$. In particular, $\gamma'' > 0$ if $z > \frac{2(m p \alpha' - (\alpha - \alpha')d)}{p \alpha m}$; hence, if $z \geq \max\left(\frac{2(p \alpha' - (\alpha - \alpha')d)}{p \alpha}, \frac{2(2\alpha' - \alpha)}{\alpha}\right)$, we get $\gamma'' > 0$.

Then, $P_n^{(3)} \leq \lambda P_n^{(2)}$ by union bound. \square

Lemma 4. *Let us denote by $P_{\text{misranking}}$ the probability that in at least one iteration, there is at least one misranking of two individuals. Then, if $z > \max\left(\frac{2(p \alpha' - (\alpha - \alpha')d)}{p \alpha}, \frac{2(2\alpha' - \alpha)}{\alpha}\right)$ and Y is large enough, $P_{\text{misranking}} \leq \delta$.*

This lemma implies that with probability at least $1 - \delta$, provided that Y has been chosen large enough, we get the same rankings of points as in the noise free case. In the noise free case Eqs. 8 and 9 hold with probability at least $1 - \delta$ - this proves the convergence with probability at least $(1 - \delta)^2$, hence the expected result; the proof of the theorem is complete. \square

Proof. (of the lemma)

We consider the probability $P_n^{(4)}$ that two individuals x_{n,i_1} and x_{n,i_2} at iteration n are misranked due to noise, so

$$\|x_{n,i_1}\|^p \leq \|x_{n,i_2}\|^p \quad (11)$$

$$\text{and } \frac{1}{Y} \sum_{j=1}^Y f(x_{n,i_1}, \omega_j) \geq \frac{1}{Y} \sum_{j=1}^Y f(x_{n,i_2}, \omega_j) \quad (12)$$

Eqs. 11 and 12 occur simultaneously if either two points have very similar fitness (difference less than δ_n) or the noise is big (larger than $\frac{\delta_n}{2}$). Therefore, $P_n^{(4)} \leq P_n^{(1)} + P_n^{(3)} \leq \lambda^2 P_n^{(0)} + \lambda P_n^{(2)} \leq (B + B') \lambda^2 \exp(-\min(\gamma', \gamma'')n)$.

$P_{\text{misranking}}$ is upper bounded by $\sum_{n \geq 1} P_n^{(4)} < \delta$ if γ' and γ'' are positive and constants large enough. γ' and γ'' can be chosen positive simultaneously if $z > \max\left(\frac{2(p \alpha' - (\alpha - \alpha')d)}{p \alpha}, \frac{2(2\alpha' - \alpha)}{\alpha}\right)$. \square

3 Experiments : how to choose the right number of resampling ?

We consider in our experiments a version of multi-membered evolution strategies, the (μ, λ) -ES, where μ denotes the number of parents and λ the number of offspring ($\mu \leq \lambda$; see Alg. 2). We denote (x_n^1, \dots, x_n^μ) the μ parents at iteration n and $(\sigma_n^1, \dots, \sigma_n^\mu)$ their corresponding step-size. At each iteration, a (μ, λ) -ES noisy algorithm : (i) generates λ offspring by mutation on the μ parents, using the corresponding mutated step-size, (ii) selects the μ best offspring by ranking

the noisy fitness values of the individuals. Thus, the current approximation of the optimum x^* at iteration n is x_n^1 , to be consistent with the previous notations, we denote $x_n = x_n^1$ and $\sigma_n = \sigma_n^1$.

Algorithm 2 An evolution strategy, with constant number of resamplings. If we consider $Y = 1$, we obtain the case without resampling. \mathcal{N}_k is a k -dimensional standard normal random variable.

Parameters : $Y > 0$, $\lambda \geq \mu > 0$, a dimension $d > 0$.

Input : μ initial points $x_1^1, \dots, x_1^\mu \in \mathbb{R}^d$ and initial step size $\sigma_1^1 > 0, \dots, \sigma_1^\mu > 0$.

$n \leftarrow 1$

while (true) **do**

 Generate λ individuals independently using :

$$\begin{aligned}\sigma_j &= \sigma_n^{\text{mod}(j-1, \mu)+1} \times \exp\left(\frac{1}{2d} \times \mathcal{N}_1\right) \\ i_j &= x_n^{\text{mod}(j-1, \mu)+1} + \sigma_j \mathcal{N}_d\end{aligned}$$

$\forall j \in \{1, \dots, \lambda\}$, evaluate i_j Y times. Let y_j be the averaging over these Y evaluations.

 Define j_1, \dots, j_λ so that $y_{j_1} \leq y_{j_2} \leq \dots \leq y_{j_\lambda}$.

Update : compute σ_{n+1}^k and x_{n+1}^k for $k \in \{1, \dots, \mu\}$:

$$\begin{aligned}\sigma_{n+1}^k &= \sigma_{j_k} \\ x_{n+1}^k &= x_{j_k}\end{aligned}$$

$n \leftarrow n + 1$

end while

Experiments are performed on the fitness function $f(x, \omega) = \|x\|^p + \|x\|^{pz/2} \mathcal{N}$, with $x \in \mathbb{R}^{15}$, $p = 2$, $z = 2.1$, $\lambda = 4$, $\mu = 2$, and \mathcal{N} a standard gaussian random variable, using a budget of 500000 evaluations. The results presented here are the mean and the median over 50 runs. The positive results are proved, above, for a given quantile of the results. This explains the good performance in Fig. 1 (median result) as soon as the number of resamplings is enough. The median performance is optimal with just 12 resamplings. On the other hand, Fig. 2 shows the mean performance of Alg. 2 with various numbers of resamplings. We see that a limited number of runs diverge so that the mean results are bad even with 16 resamplings; results are optimal (on average) for 20 resamplings.

Results are safer with 20 resamplings (for the mean), but faster (for the median) with a smaller number of resamplings.

4 Conclusion

We have shown that applying evolution strategies with a finite number of resamplings when the noise in the function decreases quickly enough near the optimum




Fig. 1: Convergence of Self-Adaptive Evolution Strategies: Median results.




Fig. 2: Convergence of Self-Adaptive Evolution Strategies: Mean results.

provides a convergence rate as fast as in the noise-free case. More specifically, if the noise decreases slightly faster than in the multiplicative model of noise, using a constant number of revaluation leads to a log-linear convergence of the algorithm. The limit case of a multiplicative noise has been analyzed in [14]; a fixed number of resamplings is not sufficient for convergence when the noise is unbounded.

Further work. We did not provide any hint for choosing the number of resamplings. Proofs based on Bernstein races [13] might be used for adaptively choosing the number of resamplings.

Acknowledgements

This paper was written during a stay in Ailab, Dong Hwa University, Hualien, Taiwan.

References

1. D. Arnold and H.-G. Beyer. Investigation of the (μ, λ) -es in the presence of noise. In *Proc. of the IEEE Conference on Evolutionary Computation (CEC 2001)*, pages 332–339. IEEE, 2001.
2. D. Arnold and H.-G. Beyer. Local performance of the $(1 + 1)$ -es in a noisy environment. *Evolutionary Computation, IEEE Transactions on*, 6(1):30–41, feb 2002.
3. D. V. Arnold and H.-G. Beyer. A general noise model and its effects on evolution strategy performance. *IEEE Transactions on Evolutionary Computation*, 10(4):380–391, 2006.
4. S. Astete-Morales, J. Liu, and O. Teytaud. log-log convergence for noisy optimization. In *Proceedings of EA 2013, LLNCS*, page accepted. Springer, 2013.
5. A. Auger. Convergence results for $(1, \lambda)$ -SA-ES using the theory of φ -irreducible Markov chains. *Theoretical Computer Science*, 334(1-3):35–69, 2005.
6. A. Auger. Linear convergence on positively homogeneous functions of a comparison-based step-size adaptive randomized search: the $(1+1)$ -es with generalized one-fifth success rule. *submitted*, 2013.
7. A. Auger, M. Jebalia, and O. Teytaud. (x, σ, η) : quasi-random mutations for evolution strategies. In *EA*, page 12p., 2005.
8. H.-G. Beyer. *The Theory of Evolution Strategies*. Natural Computing Series. Springer, Heidelberg, 2001.
9. H. Chen. Lower rate of convergence for locating the maximum of a function. *Annals of statistics*, 16:1330–1334, Sept. 1988.
10. R. Coulom. Clop: Confident local optimization for noisy black-box parameter tuning. In *Advances in Computer Games*, pages 146–157. Springer Berlin Heidelberg, 2012.
11. V. Fabian. Stochastic Approximation of Minima with Improved Asymptotic Speed. *Annals of Mathematical statistics*, 38:191–200, 1967.
12. V. Fabian. *Stochastic Approximation*. SLP. Department of Statistics and Probability, Michigan State University, 1971.

13. V. Heidrich-Meisner and C. Igel. Hoeffding and bernstein races for selecting policies in evolutionary direct policy search. In *ICML '09: Proceedings of the 26th Annual International Conference on Machine Learning*, pages 401–408, New York, NY, USA, 2009. ACM.
14. M. Jebalia, A. Auger, and N. Hansen. Log linear convergence and divergence of the scale-invariant (1+1)-ES in noisy environments. *Algorithmica*, 2010.
15. I. Rechenberg. *Evolutionstrategie: Optimierung Technischer Systeme nach Prinzipien des Biologischen Evolution*. Fromman-Holzboog Verlag, Stuttgart, 1973.
16. O. Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. *CoRR*, abs/1209.2388, 2012.
17. O. Teytaud and J. Decock. Noisy Optimization Complexity. In *FOGA - Foundations of Genetic Algorithms XII - 2013*, Adelaide, Australie, Feb. 2013.
18. O. Teytaud and H. Fournier. Lower bounds for evolution strategies using vc-dimension. In G. Rudolph, T. Jansen, S. M. Lucas, C. Poloni, and N. Beume, editors, *PPSN*, volume 5199 of *Lecture Notes in Computer Science*, pages 102–111. Springer, 2008.