

On Low Complexity Acceleration Techniques for Randomized Optimization

Sebastian Urban Stich

Institute of Theoretical Computer Science
ETH Zürich, 8092 Zürich, Switzerland
sstich@inf.ethz.ch

Abstract. Recently it was shown by Nesterov (2011) that techniques from convex optimization can be used to successfully accelerate simple derivative-free randomized optimization methods. The appeal of those schemes lies in their low complexity, which is only $\Theta(n)$ per iteration—compared to $\Theta(n^2)$ for algorithms storing second-order information or covariance matrices. From a high-level point of view, those accelerated schemes employ correlations between successive iterates—a concept looking similar to the evolution path used in Covariance Matrix Adaptation Evolution Strategies (CMA-ES). In this contribution, we (i) implement and empirically test a simple accelerated random search scheme (SARP). Our study is the first to provide numerical evidence that SARP can effectively be implemented with adaptive step size control and does not require access to gradient or advanced line search oracles. We (ii) try to empirically verify the supposed analogy between the evolution path and SARP. We propose an algorithm CMA-EP that uses only the evolution path to bias the search. This algorithm can be generalized to a family of low memory schemes, with complexity $\Theta(mn)$ per iteration, following a recent approach by Loshchilov (2014). The study shows that the performance of CMA-EP heavily depends on the spectra of the objective function and thus it cannot accelerate as consistently as SARP.

Keywords: Gradient-free optimization, accelerated random search, evolution path, adaptive step size, Covariance Matrix Adaptation, spectra.

1 Introduction

The Gradient Method [1, 2]—one of the most fundamental schemes in convex optimization—has iteration complexity $\Theta(n)$, where n is the dimension. On strongly convex functions its convergence rate is linear, depending only on the condition number of the objective function. To overcome the difficulty imposed by ill-conditioned problems, second-order methods like Newton’s method or first order Quasi-Newton methods such as the BFGS scheme [3–5] are a welcome alternative. Those schemes maintain a quadratic model of the objective function and their complexity is bounded by $\Omega(n^2)$. Limited memory schemes like L-BFGS [6, 7] trade-off linear iteration complexity $\Theta(mn)$ (where m is a fixed parameter), versus convergence rate. Accelerated versions of the Gradient Method

have linear complexity $\Theta(n)$ per iteration and converge with optimal rate among all first-order methods. On strongly convex problems the convergence rate is proportional to the square root of the condition number [1, 2, 8–10].

Randomized (gradient-free) schemes do not require first-order information, they operate by only querying function values. Such schemes are nowadays a ubiquitous tool for solving many practical problems in science and engineering where first-order information is difficult to compute or does not exist. Among the first proposed schemes that are still of considerable (theoretical) importance are Adaptive Step Size Random Search (aSS) [11] and the (almost identical) well-known (1+1)-Evolution Strategy (ES) [12] in Evolutionary Computation. More recent schemes comprise Random Pursuit (RP) [13, 14], or Random Gradient Descent [15]. Those schemes can be viewed as generalizations of the Gradient Method to zeroth-order, with iteration complexity $\Theta(n)$. Likewise, analogues of the second-order schemes try to estimate an approximation of the Hessian by finite difference computations [16, 17] or by estimating correlations among search directions. A very popular algorithm of this kind is the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [18, 19]. Limited memory variants have been proposed in [20, 21], with iteration complexity $\Theta(mn)$. Especially, the later variant due to Loshchilov shows excellent convergence in high dimensions also for small values of m . Zeroth-order analogues of the accelerated gradient schemes have been introduced [15, 22]. Those schemes massively outperform the simple random search schemes on convex problems. This performance gain does not come for free, as those schemes require valid bounds on the condition number as input parameters. However, their low iteration complexity of $\Theta(n)$ could make them a promising choice for large scale problems, where the fully-quadratic schemes inherently fail. We focus on a very simple accelerated random search scheme, which we call SARP.

By inspecting closely the accelerated search schemes, one could conclude that the difference to the classical schemes can be explained by an additional “drift” term [2, 10] that takes into account correlations of the last iterates. In the popular CMA-ES, correlations between successive iterates are accumulated in the evolution path [23]. In this work, we are interested, if the evolution path can be used for acceleration, competitive to the accelerated zeroth-order schemes for convex optimization. To this end, we introduce a variant of CMA-ES, called EP-CMA, that only uses the information stored in the evolution path to bias the direction of the search. Similar to the approach proposed in [21], this scheme can be generalized to a family of schemes, which we call EP-CMA- m . However, we do not present an efficient, i.e. $\Theta(mn)$, implementation of those schemes as this was not required here with dimension $n \leq 100$ in our empirical study.

The remainder of this paper is structured as follows. In Section 2 we present the accelerated random search scheme SARP and detail the EP-CMA- m schemes. In Section 3 we empirically test the performance of all schemes on three quadratic and the non-convex Rosenbrock function, and highlight the key results. We discuss these results and conclude the paper in Section 4.

lineSearch($\mathbf{x}, \mathbf{u}, [\sigma, p]$)	aSS($\mathbf{x}, \mathbf{u}, \sigma, p$) (<i>adaptive step size</i>)
1 if <i>exact</i> then return exactLS($\mathbf{x}, \mathbf{u}/\ \mathbf{u}\ $)	1 if $f(\mathbf{x} + \sigma\mathbf{u}) \leq f(\mathbf{x})$ then
2 else return aSS($\mathbf{x}, \mathbf{u}, \sigma, p$)	2 $\mathbf{x}_+ \leftarrow \mathbf{x} + \sigma\mathbf{u}; \sigma_+ \leftarrow \sigma \cdot \exp(1/3)$
	else
exactLS($\mathbf{x}, \mathbf{u}, [\sigma]$)	3 $\mathbf{x}_+ \leftarrow \mathbf{x}; \sigma_+ \leftarrow \sigma \cdot \exp\left(-\frac{p}{3(1-p)}\right)$
1 $\sigma_+ \leftarrow \min_{\lambda} f(\mathbf{x} + \lambda\mathbf{u}); \mathbf{x}_+ \leftarrow \mathbf{x} + \sigma_+\mathbf{u}$	4 return (\mathbf{x}_+, σ_+)
2 return (\mathbf{x}_+, σ_+)	

Fig. 1. Line search oracles for gradient-free optimization

2 Algorithms

We here present the optimization schemes considered in this study. We first detail two Random Pursuit algorithms and a simplistic variant of a standard (1+1)-CMA-ES. Then we introduce the new EP-CMA- m schemes.

RP. Random Pursuit is a basic optimization scheme that iteratively generates a sequence of approximate solutions to the global optimization problem $\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$. In each step a search direction is drawn $\mathbf{u}_k \sim \mathcal{N}(0, \mathbf{I}_n)$. In Random Pursuit with exact line search (RP-exact), first proposed in [13] and analyzed in [14], the step size σ is determined by minimizing the objective function in direction \mathbf{u} , i.e. $\sigma = \arg \min_{\lambda} f(\mathbf{x} + \lambda\mathbf{u})$. For quadratic functions $f(\mathbf{x}) := \frac{1}{2}\mathbf{x}^T \mathbf{A} \mathbf{x}$ with Hessian A , the expected one-step progress can be estimated as:

$$\mathbb{E}[f(\mathbf{x}_+) \mid \mathbf{x}] \leq (1 - \lambda_{\min}(A)/\text{Tr}[A]) f(\mathbf{x}), \quad (1)$$

where \mathbf{x} is the current iterate, $\mathbf{x}_+ := \mathbf{x} + \sigma\mathbf{u}$ denotes the next iterate. This statement can also be generalized to arbitrary smooth convex functions [14]. Stich et al. [14] show that RP-exact still converges if the line search is not performed exactly, but allowing relative errors. Therefore, we also consider Random Pursuit with adaptive step sizes (RP) instead of exact line search. In RP the step size is dynamically controlled such as to approximately guarantee a certain probability p of finding an improving iterate. Depending on the underlying test function, different optimality conditions can be formulated for the value p . Schumer and Steiglitz [11] suggest the setting $p = 0.27$ which is considered throughout this work. We use immediate exponential step size control as explicitly formulated in the aSS sub-routine in Fig. 1. RP is identical to the well known (1+1)-ES.

SARP. Accelerated random search schemes are fundamentally different from the simple random search schemes. Instead of generating only one sequence of iterates, those algorithms typically maintain two or more sequences simultaneously (here essentially \mathbf{x}_k and \mathbf{y}_k , see Fig. 2). Those sequences allow to store gathered knowledge on the objective function which yields better performance. In Fig. 2 we present a simple version of the accelerated random search scheme proposed in [14] and refer to it as simple accelerated random search (SARP). Like RP, SARP can (in practice) be used with exact line search oracles or with adaptive step size control, although convergence for those oracles has not been proven yet. For Nesterov’s accelerated random search scheme [15], the expected one-step progress can be estimated as

<hr/> <p>RP($\mathbf{x}_0, N, [\sigma_0, p]$)</p> <hr/> <pre> 1 for $k = 1$ <i>to</i> N do 2 $\mathbf{u}_k \sim \mathcal{N}(0, I_n)$ 3 $(\mathbf{x}_k, \sigma_k) \leftarrow \text{lineSearch}(\mathbf{x}_{k-1}, \mathbf{u}_k, [\sigma_{k-1}])$ 4 return \mathbf{x}_N </pre> <hr/> <p>EP-CMA-m ($\mathbf{x}_0, N, \sigma_0, p, c_c, c_{\text{cov}}$)</p> <hr/> <pre> 1 $\hat{\mathbf{p}}_0 \leftarrow \mathbf{0}; \hat{\mathbf{p}}_1, \dots, \hat{\mathbf{p}}_{m-1} \leftarrow \mathbf{0}; q = 0$ 2 for $k = 1$ <i>to</i> N do 3 $C_k \leftarrow I_n$ 4 for $i = 1$ <i>to</i> $m-1$ do 5 $C_k \leftarrow (1 - c_{\text{cov}})C_k + c_{\text{cov}}\hat{\mathbf{p}}_i\hat{\mathbf{p}}_i^T$ 6 $\mathbf{u}_k \sim \mathcal{N}(0, C_k)$ 7 $(\mathbf{x}_k, \sigma_k) \leftarrow \text{aSS}(\mathbf{x}_{k-1}, \mathbf{u}_k, \sigma_{k-1})$ 8 $\mathbf{y}_k \leftarrow (\mathbf{x}_k - \mathbf{x}_{k-1})/\sigma_{k-1}$ 9 if $\mathbf{y}_k \neq \mathbf{0}$ (<i>success</i>) then 10 $\mathbf{p}_k \leftarrow (1 - c_c)\mathbf{p}_{k-1} + \sqrt{c_c(2 - c_c)}\mathbf{y}_k$ 11 else $\mathbf{p}_k \leftarrow (1 - c_p)\mathbf{p}_{k-1}$ if $k > q + n^2/m$ then 12 $\hat{\mathbf{p}}_1 \leftarrow \hat{\mathbf{p}}_2, \dots, \hat{\mathbf{p}}_{m-2} \leftarrow \hat{\mathbf{p}}_{m-1}; q = k$ 13 return \mathbf{x}_N </pre> <hr/>	<hr/> <p>SARP($\mathbf{x}_0, N, m, L, [\sigma_0]$)</p> <hr/> <pre> 1 $\mathbf{y}_0 \leftarrow \mathbf{x}_0; \mathbf{v}_0 \leftarrow \mathbf{x}_0; \theta \leftarrow \sqrt{\frac{m}{2n^2L}}$ 2 for $k = 1$ <i>to</i> N do 3 $\mathbf{u}_k \sim \mathcal{N}(0, I_n)$ 4 $(\mathbf{x}_k, \sigma_k) \leftarrow \text{lineSearch}(\mathbf{y}_{k-1}, \mathbf{u}_k, [\sigma_{k-1}])$ 5 $\mathbf{y}_k \leftarrow (\theta\mathbf{v}_{k-1} + \mathbf{x}_k)/(1 + \theta)$ 6 $\mathbf{v}_k \leftarrow (1 - \theta)\mathbf{v}_{k-1} + \theta\mathbf{y}_k + \theta n \frac{L}{m} \sigma_k \mathbf{u}_k$ 7 return \mathbf{x}_N </pre> <hr/> <p>(1+1)-CMA($\mathbf{x}_0, N, \sigma_0, p, c_c, c_{\text{cov}}$)</p> <hr/> <pre> 1 $C_0 \leftarrow I_n; \mathbf{p}_0 \leftarrow \mathbf{0}$ 2 for $k = 1$ <i>to</i> N do 3 $\mathbf{u}_k \sim \mathcal{N}(0, C_{k-1})$ 4 $(\mathbf{x}_k, \sigma_k) \leftarrow \text{aSS}(\mathbf{x}_{k-1}, \mathbf{u}_k, \sigma_{k-1})$ 5 $\mathbf{y}_k \leftarrow (\mathbf{x}_k - \mathbf{x}_{k-1})/\sigma_{k-1}$ 6 if $\mathbf{y}_k \neq \mathbf{0}$ (<i>success</i>) then 7 $\mathbf{p}_k \leftarrow (1 - c_c)\mathbf{p}_{k-1} + \sqrt{c_c(2 - c_c)}\mathbf{y}_k$ 8 $C_k \leftarrow (1 - c_{\text{cov}})C_{k-1} + c_{\text{cov}}\mathbf{p}_k\mathbf{p}_k^T$ 9 else 10 $C_k \leftarrow C_{k-1}; \mathbf{p}_k \leftarrow (1 - c_p)\mathbf{p}_{k-1}$ 10 return \mathbf{x}_N </pre> <hr/>
---	--

Fig. 2. RP, EP-CMA and CMA-ES schemes

$$\mathbb{E}[f(\mathbf{x}_+) \mid \mathbf{x}] \leq (1 - (n\sqrt{\kappa})^{-1}) f(\mathbf{x}), \quad (2)$$

where condition $\kappa = L/m$ and the two parameters $m \leq \lambda_{\min}(A)$ and $L \geq \lambda_{\max}(A)$ are required as input to the algorithm (and always provided in our numerical study). This rate is much better than (1) and we hope to see that SARP attains comparable performance. SARP is not a monotone scheme, that is, the function values of the iterates are not monotonically decreasing. SARP is closely related to the first-order accelerated search scheme of Nesterov [2]. This scheme also simultaneously maintains two sequences \mathbf{x}'_k and \mathbf{y}'_k of iterates (but requires access to the gradient in every iteration). For Nesterov's first-order scheme it is known [2, p.79] that the sequence \mathbf{y}'_k obeys

$$\mathbf{y}'_{k+1} = \mathbf{x}'_{k+1} + \beta' (\mathbf{x}'_{k+1} - \mathbf{x}'_k), \quad (3)$$

for $\beta' = 1 - 2/\sqrt{\kappa} + O(1/\kappa)$. Thus the additional $(\mathbf{x}'_{k+1} - \mathbf{x}'_k)$ acts like a drift term, cf. [1]. For SARP with parameter $\theta' = \sqrt{1/(n^2\kappa)}$ (only slightly different from θ in Fig. 2) the same reformulation of the update reveals

$$\mathbf{y}_{k+1} = \mathbf{x}_k + \beta (\mathbf{x}_{k+1} - \mathbf{x}_k), \quad (4)$$

for $\beta = (1 - \theta')/(1 + \theta') = 1 - 2/(n\sqrt{\kappa}) + O(1/\kappa)$. The main term contributing to the drift is approximately only an $1/n$ -fraction of the step, accounting for the uncertainty emerging from the randomness.

(1+1)-CMA-ES. In contrast to the presented Random Pursuit schemes, in CMA-ES new search points are sampled from a multivariate normal distribution

$\mathbf{u}_k \sim \mathcal{N}(0, C_k)$ whose parameter C_k is updated in each iteration based on the evaluation of the samples. The covariance matrix can be adapted using different rank-1 [18, 24] or rank-k updates [19]. In addition, the CMA-ES scheme is augmented by an auxiliary variable called evolution path that takes into account the correlation of successive means taken over a finite horizon. In [18, 23], the evolution path \mathbf{p}_k is updated as

$$\mathbf{p}_{k+1} = (1 - c_c)\mathbf{p}_k + \sqrt{c_c(1 - c_c)}\mathbf{u}_k. \quad (5)$$

Cumulative information about successive steps is stored in the variable \mathbf{p}_k . We use a simplistic CMA-ES variant, closely following [18], see Fig. 2. We use the simple Adaptive Step Size control aSS to determine the step size σ_k , the covariance matrix update solely uses the information of the evolution path like in [24] and for simplicity we refrain from implementing any regularization features, in contrast to [24]. We use the same parameters that were proposed in [24] for the (1+1)-CMA-ES, namely $c_c = 2/(n + 2)$, $c_p = 1/12$ and $c_{\text{cov}} = 2/(n^2 + 6)$.

EP-CMA-1. The evolution path \mathbf{p}_k accumulates information over successful steps. This accumulation can be seen as a smoothing of the noisy information obtained in single steps, at the effect that the evolution path points into direction of more promising function values [18]. In this study, we are interested if the evolution path can be used in a similar way as the drift term in (3) or (4), respectively, to accelerate the search. There are several ways to incorporate the evolution path \mathbf{p}_k into the update scheme. We suggest to use the path \mathbf{p}_k in the following way: in the simple random search scheme RP (equivalent to (1+1)-ES), we sample in iteration k a direction from $\mathbf{u}_k \sim \mathcal{N}(0, (1 - c_{\text{cov}})I_n + c_{\text{cov}}\mathbf{p}_k\mathbf{p}_k^T)$, with bias along the direction indicated by \mathbf{p}_k . This has the effect that we only follow successful steps, but the drift imposed by the evolution path might be smaller than it ideally should be. The scheme EP-CMA-1 is detailed in Fig 2, we used c_c and c_p as above, and $c_{\text{cov}} = 1/5$. This approach is similar to [25].

EP-CMA-m. The proposed EP-CMA-1 can easily be generalized to a whole family of optimization schemes by an approach presented in [21]. In EP-CMA-1, only the information stored in the current evolution path \mathbf{p}_k is used to bias the search direction. But we could also afford to temporarily store a small number m of past $\mathbf{p}_{k'}$ for $k' < k$, and use the information collectively to bias the search. As two successive evolution paths are likely highly correlated, we propose to store the evolution path only every n^2/m -th generation (and up to at most $(m - 1)$ copies simultaneously). The resulting scheme is detailed in Fig. 2. We used $c_c = 2/(n + 2)$ as in CMA-ES, and for $m > 1$, $c_{\text{cov}} = 2/(6 + m)$ for EP-CMA- m . If implemented carefully, EP-CMA- m has $\Theta(mn)$ complexity per iteration (not shown in Fig 2). For $m = n^2$, the updates of EP-CMA- m are identical to the updates of (1+1)-CMA-ES, if *limited* to a finite horizon of n^2 steps. In contrast, the low memory method proposed in [20] behaves similar to CMA-ES already for $m = n$, but has iteration complexity $\Theta(nm^2)$.

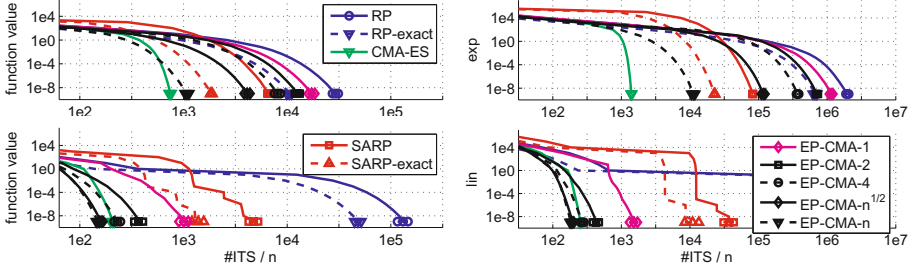


Fig. 3. Evolution of FVAL vs. #ITS on f_{exp} (top) and f_{lin} (bottom) with $L = 1\text{E}4$ (left) and $L = 1\text{E}6$ (right) in $n = 100$ dimensions. For 51 (11 for RP on f_{lin} with $L = 1\text{E}6$) runs we recorded #ITS needed to reach FVAL of $1\text{E}-9$. The trajectory realizing the median values is depicted, mean and one standard deviation are indicated by markers. (RP on f_{lin} with $L = 1\text{E}6$ reaching FVAL $< 1\text{E}-2$ after $1\text{E}6.5n$ #ITS.)

3 Empirical Study

We now present the setup of our empirical study. We focus on the following schemes: (i) the two Random Pursuit schemes with adaptive step size control (denoted as RP and SARP) and with exact line search (denoted as RP-exact and SARP-exact), (ii) the simplified (1+1)-CMA-ES and (iii) the EP-CMA- m schemes as introduced in Sec. 2, see Fig. 2. We use EP-CMA- m with parameters $m = 1, 2, 4, \sqrt{n}, n$. This totals in 10 different schemes, all of which were implemented in MATLAB and will be made available on the authors website.

We tested the performance of all algorithms on three variants of the ellipsoidal benchmark function [18] and the non-convex Rosenbrock function, detailed in Table 1. The quadratic functions were chosen in such a way that the extremal values of their spectra (1 and L) both agree. We considered the quadratic functions with parameters $L = 1\text{E}4$ and $L = 1\text{E}6$ each, and repeat the experiments in dimensions $n = 20, 40, 60, 80, 100$.

For all experiments, initial settings were $\mathbf{x}_0 = \mathbf{1}$, $\sigma_0 = 1$ and $p = 0.27$ (for schemes with the aSS routine). We count the number of iterations (#ITS) needed to decrease the function value (FVAL) below $1\text{E}-9$. A graphical summary of our results can be found in Fig. 3-5. Results not depicted here are reported in the supplementary online material [26]. We now proceed by discussing some of the key results.

Table 1. List of benchmark functions

$f_{\text{exp}}(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^n L^{\frac{i-1}{n-1}} x_i^2$	$f_{\text{rosen}}(\mathbf{x}) = \sum_{i=1}^{n-1} \left(100 \cdot (x_i^2 - x_{i+1})^2 + (x_i - 1)^2 \right)$
$f_{\text{lin}}(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^n \left(1 + i \frac{(L-1)}{(n-1)} \right) x_i^2$	$f_{\text{two}}(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^{\lfloor n/2 \rfloor} x_i^2 + \frac{L}{2} \sum_{i=\lfloor n/2 \rfloor}^n x_i^2$




Fig. 4. Evolution of FVAL vs. #ITS on f_{exp} with $L = 1\text{E}4$ (left) and $L = 1\text{E}6$ (right) in $n = 20$ and $n = 80$ dimensions. For 51 runs we recorded #ITS needed to reach FVAL of $1\text{E}-9$. The trajectory realizing the median values is depicted, mean and one standard deviation are indicated by markers.

Line Search. Both RP and SARP were tested with exact line search oracle and adaptive step size control. In Fig. 3 we see that the exact schemes outperform their adaptive variants in $n = 100$ dimensions by a factor of roughly 2-3. This pattern is observed throughout the whole benchmark in all dimensions. Thus we omit to display the results for exact line search in subsequent Figs. 4-5.

SARP vs. EP-CMA-1. The picture is twofold. In Fig. 3 we see that EP-CMA-1 outperforms SARP by a factor of roughly 5 on f_{lin} with $L = 1\text{E}4$ (factor 24 for $L = 1\text{E}6$). The smallest eigenvalue of this function is separated from the second largest by a gap of roughly n . Hence, knowledge of one important direction reduces the conditioning of the function by a large factor. This factor becomes smaller in higher dimension. This scaling in the dimensions is indeed observed empirically, and depicted in [26].

On the other three functions, SARP performs consistently better than EP-CMA-1. On f_{exp} with $L = 1\text{E}4$ in $n = 100$ dimensions the factor is roughly 3, its roughly 14 for $L = 1\text{E}6$ (Fig. 3), and exceeds 10 on both f_{two} and f_{rosen} (Fig. 5). Considering the scaling in dimension (Figs. 4-5; and [26]), we observe that the relative performance ($\#ITS/n$) of SARP remains constant on all four benchmark functions, as predicted by theory for a similar method [15, 22].

EP-CMA-schemes. The EP-CMA- m schemes consistently work better for increasing values of m throughout the whole benchmark (Figs. 3-5). On f_{exp} with $L = 1\text{E}4$ the difference in #ITS between EP-CMA- n and EP-CMA-1 is roughly a factor of 10, and 20 for $L = 1\text{E}6$ (Fig. 3). The gap becomes gradually smaller on f_{lin} , f_{two} (especially for $L = 1\text{E}6$, see [26]), and is insignificant on f_{rosen} (Fig. 5). On f_{lin} the EP-CMA- m schemes perform extremely well, already for small m . EP-CMA-4 performs approximately as good as CMA-ES, for both parameters $L = 1\text{E}4$ and $L = 1\text{E}6$ (Fig. 3). On f_{exp} in $n = 100$ dimensions and parameter $L = 1\text{E}4$, both SARP and EP-CMA-4 need about the same #ITS. For parameter $L = 1\text{E}4$ the performance of SARP is the same as the performance of EP-CMA- \sqrt{n} (Fig. 3). On both f_{two} and f_{rosen} , the EP-CMA- m scheme cannot reach the performance of SARP, though on f_{rosen} the EP-CMA- m schemes perform as good as CMA-ES (Fig. 5).




Fig. 5. Evolution of FVAL vs. #ITS on f_{two} with $L = 1\text{E}4$ (left) and f_{rosen} (right) in $n = 20$ and $n = 80$ dimensions. For 51 runs we recorded #ITS needed to reach FVAL of $1\text{E}-9$. The trajectory realizing the median values is depicted, mean and one standard deviation are indicated by markers.

CMA-ES. Fig. 4 shows nicely the quadratic dependence of the performance of CMA-ES on the dimension n , see also [26] where we report the data for all considered dimensions. The #ITS of the Random Pursuit schemes (RP, SARP) to reach the target accuracy increases only linearly (the relative performance ($\#ITS/n$) is constant over the dimensions). In the dimensions $n \leq 100$ considered here, CMA-ES is the best performing scheme on f_{exp} (Fig. 3) and f_{two} (Fig. 5); on f_{lin} the EP-CMA- m schemes match its performance for $m \geq 4$ (Fig. 3). A notable exception is the behavior on the non-convex f_{rosen} , where only SARP can accelerate and the other schemes, including CMA-ES, require over 10 times more #ITS to reach the same accuracy (Fig. 5).

4 Discussion and Conclusions

In this contribution we emphasize the importance of accelerated random gradient schemes [15, 22]. Each iteration in SARP has only linear complexity, yet the scheme takes correlations between successive iterates into account. In CMA-ES, such correlations are collected in the evolution path [18, 23] and stored in the covariance matrix. This requires $\Theta(n^2)$ simple operations per iteration. The proposed EP-CMA-1 uses as well the information of the evolution path to bias the search, but does not store a full-rank covariance matrix.

Line Search. We empirically tested two Random Pursuit algorithms with an exact line search oracle. Such an oracle is in general not available for general black-box optimization problems and the line search must for instance be implemented as bisection search (cf. [14, 27]) at the expense of additional function evaluations per iteration. The empirical data shows that both Random Pursuit algorithms do perform well if a simple adaptive step size scheme is used instead of the line search. This makes both schemes (especially SARP) promising candidates for black-box optimization, also in high dimension, as the runtime scales only linearly with the dimension. Up to our knowledge, no experimental results for SARP with adaptive step size have been published yet (the authors in [14] considered a line search with high accuracy, almost like SARP-exact).

Acceleration in EP Schemes. Our empirical results show that the sole use of the evolution path can lead to astonishingly good performance—depending on the problem and its eigenvalue spectrum. The speed-up of EP-CMA-1 on f_{lin} can be explained by the fact that the condition number of the problem drops once the algorithm has learned the most insensitive direction. Hence, the acceleration can be explained by formula (1) rather than (2). For SARP the situation is more promising. The data indicates that the convergence on f_{exp} and f_{two} is as described in (2). The same seems to be true on the non-convex f_{rosen} where SARP needs an order of magnitude less #ITS than all other schemes, including CMA-ES. Only on f_{lin} this does not to hold, as SARP is only one order of magnitude faster than RP. Consider the update (4). By expansion we obtain

$$\mathbf{y}_{k+1} = \mathbf{x}_k + \beta(\sigma_{k+1}\mathbf{u}_{k+1} + \beta(\mathbf{x}_k - \mathbf{x}_{k-1})) = \mathbf{x}_k + \sum_{i=1}^{k+1} \beta^{k+2-i} \sigma_i \mathbf{u}_i. \quad (6)$$

We see that the drift is a weighted average of the previous steps $\sigma_i \mathbf{u}_i$. The discount factor β is the expected convergence rate. Therefore, the influence of a step $\sigma_i \mathbf{u}_i$ on \mathbf{y}_{k+1} is roughly the same for all $i = 1, \dots, k$. In contrast to this, the evolution path \mathbf{p}_k stores only information of the directions of the last steps (but no step sizes). The discount factor is approximately $1 - 2/n$. Although the evolution path \mathbf{p}_k is a cumulation of all old steps, the weight of old steps is exponentially small compared to the influence of the newest steps. We might conclude that the mechanism of accelerated random schemes like SARP is therefore inherently different to the concept of the evolution path, supporting reports in [27]. However, we cannot rule out the possibility, that with a different choice of internal parameters of EP-CMA-1 the difference to SARP could be reduced.

Limited Memory Schemes. The performance of the proposed EP-CMA- m schemes uniformly increases for larger parameters m , as well as the complexity of each single iteration. An optimal trade-off for the parameter m has to be found, depending on the dimension n and the cost of individual function evaluations. The data shows that the EP-CMA- m schemes can dramatically improve the performance of simple random search already for small values of m . The speed-up depends crucially on the eigenvalue spectra of the objective function. It seems that these schemes can not reach the performance of the related variants in [21].

We generally conclude, that the here proposed algorithmic schemes with linear iteration complexity could be a promising way to handle high dimensional black-box optimization problems. However, the empirical data suggest that there is an intrinsic limitation for the EP schemes, as they depend on the eigenvalue spectrum of the objective function. This behavior is not observed for SARP. We like to advocate that features of accelerated schemes (like SARP) should therefore be taken seriously into account when facing high dimensional problems.

References

1. Polyak, B.: Introduction to Optimization. Optimization Software - Inc. (1987)
2. Nesterov, Y.: Introductory Lectures on Convex Optimization. Kluwer (2004)

3. Broyden, C.G.: The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations. *IMA J. of Appl. Math.* 6(1), 76–90 (1970)
4. Fletcher, R.: A new approach to variable metric algorithms. *The Computer Journal* 13(3), 317–322 (1970)
5. Goldfarb, D.: A Family of Variable-Metric Methods Derived by Variational Means. *Mathematics of Computation* 24(109), 23–26 (1970)
6. Nocedal, J.: Updating Quasi-Newton Matrices with Limited Storage. *Mathematics of Computation* 35(151), 773–782 (1980)
7. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. *Mathematical Programming* 45(1-3), 503–528 (1989)
8. Nesterov, Y.: A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady* 27(2), 372–376 (1983)
9. Nesterov, Y.: Smoothing technique and its applications in semidefinite optimization. *Mathematical Programming* 110(2), 245–259 (2007)
10. Tseng, P.: On accelerated proximal gradient methods for convex-concave optimization. Submitted to *SIAM Journal on Optimization* (2008)
11. Schumer, M., Steiglitz, K.: Adaptive step size random search. *IEEE Transactions on Automatic Control* 13(3), 270–276 (1968)
12. Rechenberg, I.: Evolutionsstrategie; Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Frommann-Holzboog (1973)
13. Mutseniyeys, V.A., Rastrigin, L.A.: Extremal control of continuous multi-parameter systems by the method of random search. *Eng.Cyb.* 1, 82–90 (1964)
14. Stich, S.U., Müller, C.L., Gärtner, B.: Optimization of convex functions with Random Pursuit. *SIAM Journal on Optimization* 23(2), 1284–1309 (2013)
15. Nesterov, Y.: Random Gradient-Free Minimization of Convex Functions. Technical report, ECORE (2011)
16. Leventhal, D., Lewis, A.S.: Randomized Hessian estimation and directional search. *Optimization* 60(3), 329–345 (2011)
17. Stich, S.U., Gärtner, B., Müller, C.L.: Variable Metric Random Pursuit (2012) (submitted), <http://arxiv.org/abs/1210.5114>
18. Hansen, N., Ostermeier, A.: Completely Derandomized Self-Adaption in Evolution Strategies. *Evolutionary Computation* 9(2), 159–195 (2001)
19. Hansen, N., Muller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evol. Comput.* 11(1), 1–18 (2003)
20. Knight, J.N., Lunacek, M.: Reducing the Space-time Complexity of the CMA-ES. In: *GECCO 2007*, pp. 658–665. ACM (2007)
21. Loshchilov, I.: A Computationally Efficient Limited Memory CMA-ES for Large Scale Optimization. To appear *GECCO* (2014), <http://arxiv.org/abs/1404.5520>
22. Lee, Y.T., Sidford, A.: Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems. In: *FOCS*, pp. 147–156. IEEE (2013)
23. Ostermeier, A., Gawelczyk, A., Hansen, N.: Step-size adaptation based on non-local use of selection information. In: Davidor, Y., Männer, R., Schwefel, H.-P. (eds.) *PPSN 1994. LNCS*, vol. 866, pp. 189–198. Springer, Heidelberg (1994)
24. Igel, C., Suttorp, T., Hansen, N.: A Computational Efficient Covariance Matrix Update and a (1+1)-CMA for Evolution Strategies. In: *GECCO*, pp. 453–460 (2006)

25. Sun, Y., Schaul, T., Gomez, F., Schmidhuber, J.: A linear time natural evolution strategy for non-separable functions. In: Proc. 15th Genetic and Evolutionary Computation Conference Companion, pp. 61–62. ACM (2013)
26. Stich, S.U.: Supplementary Online Mat (2014), <http://arxiv.org/abs/1406.2010>
27. Stich, S.U., Müller, C.L.: On Spectral Invariance of Randomized Hessian and Covariance Matrix Adaptation Schemes. In: Coello, C.A.C., Cutello, V., Deb, K., Forrest, S., Nicosia, G., Pavone, M. (eds.) PPSN 2012, Part I. LNCS, vol. 7491, pp. 448–457. Springer, Heidelberg (2012)