Evolution under Strong Noise: A Self-Adaptive Evolution Strategy Can Reach the Lower Performance Bound - the pcCMSA-ES

Michael Hellwig[1,2] and Hans-Georg Beyer[1,3]

[1] Vorarlberg University of Applied Sciences, Research Center PPE,
Hochschulstraße 1, 6850 Dornbirn, Austria.
[2] michael.hellwig@fhv.at, WWW home page: http://homepages.fhv.at/hemi
[3] hans-georg.beyer@fhv.at, WWW home page: http://homepages.fhv.at/hgb

Abstract. According to a theorem by Astete-Morales, Cauwet, and Teytaud, "simple Evolution Strategies (ES)" that optimize quadratic functions disturbed by additive Gaussian noise of constant variance can only reach a simple regret log-log convergence slope $\geq -1/2$ (lower bound). In this paper a population size controlled ES is presented that is able to perform better than the $-1/2$ limit. It is shown experimentally that the pcCMSA-ES is able to reach a slope of -1 being the theoretical lower bound of *all* comparison-based direct search algorithms.

1 Introduction

In many real-world applications the problem complexity is increased by noise. Noise can stem from different sources such as randomized simulations or sensory disturbances. Evolutionary Algorithms (EAs) proved to be successful for optimization in the presence of noise [1, 2]. However, the performance of the EAs degrades under strong noise and can even prevent the EA from converging to the optimizer.

Performance of EAs is usually measured by the amount of objective function evaluations n needed to reach a certain expected fitness compared to the non-noisy objective function value at the optimizer. This quantity is sometimes referred to as *simple regret* $\mathrm{SR}(n)$. It is defined in the case of minimization of the noisy function $\tilde{f}(\boldsymbol{y}), \boldsymbol{y} \in \mathbb{R}^N$ as

$$\mathrm{SR}(n) := \mathrm{E}[\tilde{f}(\boldsymbol{y}_n)] - f(\hat{\boldsymbol{y}}), \tag{1}$$

where the noisy fitness $\tilde{f}(\boldsymbol{y})$ is given by $\tilde{f}(\boldsymbol{y}) = f(\boldsymbol{y}) + \delta$ and $\boldsymbol{y}_n$ is the object vector recommended by the EA after n $\tilde{f}(\boldsymbol{y})$ evaluations. $f(\boldsymbol{y})$ is the deterministic objective function to be optimized which is disturbed by unbiased noise δ. The minimizer of $f(\boldsymbol{y})$ is denoted as $\hat{\boldsymbol{y}}$. The random variate δ describes the noise, which may or may not scale with the objective function value

$$\text{a): } \delta \sim \sigma_{\epsilon_r} f(\boldsymbol{y}) \mathcal{N}(0,1) \qquad \text{and} \qquad \text{b): } \delta \sim \sigma_\epsilon \mathcal{N}(0,1), \tag{2}$$

and is assumed to be normally distributed with standard deviation $\sigma_{\epsilon_r}|f(\boldsymbol{y})|$ and σ_ϵ, respectively. The quantity σ_ϵ is referred to as noise strength.

There are different options to tackle the performance degradation of EAs that can be basically subdivided into two classes:

(i) reducing the noise observed by the EA by use of *resampling*, i.e. averaging over a number of κ objective function values (for a fixed $\boldsymbol{y}$), and
(ii) handling the noise by successively increasing the population size.

However, both methods implicate an increase of the required number n of fitness evaluations. In order to avoid a unnecessary excess of function evaluations, the question arises at which point to take the countermeasures (ii) or (i), i.e. to increase the population size or to use the $\tilde{f}$-averaging. As far as option a) is concerned, there is a definite answer regarding the $(\mu/\mu_I, \lambda)$-Evolution Strategy (ES) on quadratic functions [3, 4]: It is better to increase the population size than to perform resampling.

No matter whether one uses option (i) or (ii), in both cases techniques are required to detect the presence of noise. This can be easily done by resampling a candidate solution ($\kappa = 2$) because noise is reflected in changes of a candidate solution's measured fitness of two consecutive evaluations (for fixed $\boldsymbol{y}$). However, small noise strengths are usually well tolerated by the ES. That is, the ES can still approach the optimizer. In such cases there is no need to handle this noise. Another approach introduced in the UH-CMA-ES [5] considers the rank changes within the offspring individuals after resampling the population with $\kappa = 2$. If there are no or only a few rank changes, one can assume that the noise does not severely disturb the selection process. This approach is interesting, but seems still to be too pessimistic, i.e., even if there is a lot of rank changes, there may be still progress towards the optimizer due to the genetic repair effect taking place by the intermediate recombination operator. In [4] a population size control rule was proposed which is based on the residual error. The dynamics of the $(\mu/\mu_I, \lambda)$-ES in a noisy environment with constant noise strength σ_ϵ will usually approach a steady state in a certain distance to the optimizer. At that point, fluctuations of the parental fitness values around their mean value can be observed. The population size is then increased if the fitness dynamics on average does not exhibit further progress.

This paper presents a new detection method which is based on a linear regression analysis of the noisy fitness dynamics. Estimating the slope of the linear regression line, the direction of the trend can be determined. However, the estimated slope is a random variate. Therefore, a hypothesis test must be used to check the significance of the observed trend. If there is not a significant fitness decrease tendency, the population size will be increased. In the opposite case the population size can be decreased (up to a predefined limit). This approach is integrated into the covariance matrix self-adaptation evolution strategy (CMSA-ES) [6] yielding the *population controlled* (pc)CMSA-ES.

The applicability of the proposed algorithm is demonstrated on the noisy ellipsoid model. Investigating the SR(n) performance dynamics of the pcCMSA-ES in the *strong noise* scenario $\sigma_\epsilon = \text{const.}$ (i.e., the noise does *not* vanish at the optimizer), a remarkable observation can be made: SR(n) $\approx c/n$. That is, the slope of the log-log plot reaches -1 approximately for sufficiently large number n of function evaluations. This is in contrast to a Theorem derived in [7]. There it is stated that *simple* ES can only reach a slope $\geq -1/2$, no matter whether one uses resampling or population upgrading. Remarkably, the -1 slope actually observed already represents the lower performance bound that cannot be beaten by any direct search algorithm as has been proven in [8].

The rest of this paper is organized as follows. The proposed noise detection technique by linear regression analysis is presented in Sec. 2. This technique is used to extend the CMSA-ES with a population control rule in Sec. 3 yielding the pcCMSA-ES. Empirical investigations are provided and discussed in Sec. 4. The paper closes with a summary and a outlook at future research questions.

2 Stagnation Detection by Use of Linear Regression Analysis

Stagnation or divergence behavior coincides with a non-negative trend within the observed fitness value dynamics of the ES (minimization considered). For trend analysis a regression model of the parental centroid fitness sequence of length L is used. If the slope of this model is *significantly* negative, the ES converges. In the opposite case, the population size must be increased. The decision will be based on statistical hypothesis testing.

Considering a not too long series of observed parental centroid fitness values, the observed time series can be approximated piecewise by a linear regression model. That is, a straight line is fitted through a set of L data points $\{(x_i, f_i),\ i = 1, \ldots, L\}$ in such a manner that the sum of squared residuals of the model

$$f_i = ax_i + b + \epsilon_i \tag{3}$$

is minimal. Here ϵ_i models the random fluctuations. Determining the optimal a and b is a standard task yielding [9]

$$\hat{a} = \frac{\sum_{i=1}^{L} (x_i - \bar{x})(f_i - \bar{f})}{\sum_{i=1}^{L} (x_i - \bar{x})^2} \qquad \text{and} \qquad \hat{b} = \bar{f} - \hat{a}\bar{x}, \tag{4}$$

where $\bar{x}$ and $\bar{f}$ represent the sample means of the observations. Due to the ϵ_i random fluctuations the *estimate* $\hat{a}$ itself is a random variate. Therefore, the real (but unknown) a value can only be bracketed in a confidence interval. Assuming L sufficiently large, the central limit theorem guarantees that the estimator $\hat{a}$ of a is asymptotically normally distributed with mean a. Thus, the sum of squared residuals $\sum_{i=1}^{L}(f_i - b - ax_i)^2$ is distributed proportionally to χ^2_{L-2} with $L - 2$ degrees of freedom and is independent of $\hat{a}$, cf. [9]. This allows to construct a test statistic

$$T_{L-2} = \frac{\hat{a} - a}{s_{\hat{a}}} \qquad \text{with} \qquad s_{\hat{a}} = \sqrt{\frac{\sum_{i=1}^{L}(f_i - b - ax_i)^2}{(L-2)\sum_{i=1}^{L}(x_i - \bar{x})^2}}, \tag{5}$$

where T_{L-2} is a t-distributed random variate with $L - 2$ degrees of freedom [9].

Since $\hat{a}$ is a random variate, an observed $\hat{a} < 0$ does not guarantee convergence. Therefore, a hypothesis test will be used to put the decision on a statistical basis. Let $\mathrm{H}_0 : a \geq 0$ be the hypothesis that the ES increases the population size (because of non-convergence). We will only reject H_0 if there is significant evidence for the alternative $\mathrm{H}_1 : a < 0$. (In the latter case, the population size will not be increased.) That is, a left tailed test is to be performed with a significance level α (probability of wrongly rejecting H_0), i.e. $\Pr[\hat{a} < c|\mathrm{H}_0] = \alpha$, where c is the threshold (to be determined) below which the correct H_0 is rejected with error probability α. Resolving the left equation in (5) for $\hat{a}$ yields $\hat{a} = a + s_{\hat{a}}T_{L-2}$ and therefore $\Pr[a + s_{\hat{a}}T_{L-2} < c|\mathrm{H}_0] = \alpha$. This is equivalent to $\Pr[T_{L-2} < (c-a)/s_{\hat{a}}|\mathrm{H}_0] = \alpha$. Noting that $\Pr[T_{L-2} < (c-a)/s_{\hat{a}}] = F_{T_{L-2}}((c-a)/s_{\hat{a}})$ is the cdf of T_{L-2}, one can apply the quantile function yielding $(c-a)/s_{\hat{a}} = t_{\alpha;L-2}$, where $t_{\alpha;L-2}$ is the α quantile of the t-distribution with $L-2$ degrees of freedom. Solving for c one obtains $c = a + s_{\hat{a}} t_{\alpha;L-2}$. Thus, $c \geq s_{\hat{a}} t_{\alpha;L-2}$ and as threshold ($a = 0$) one gets $c = s_{\hat{a}} t_{\alpha;L-2}$. That is, if

$$\hat{a} < s_{\hat{a}} t_{\alpha;L-2} \tag{6}$$

H_0 is rejected indicating a significant negative trend (i.e., convergence towards the optimizer, no population size increase needed).

3 The pcCMSA-ES Algorithm

Combining the convergence hypothesis test of Sec. 2 with the basic $(\mu/\mu_I, \lambda)$-CMSA-ES introduced in [6] an ES with adaptive population size control, the *population control* (pc)CMSA-ES is presented in Alg. 1. Until the algorithm has generated a list $\mathcal{F}$ of L parental centroid function values an ordinary CMSA-ES run with truncation ratio ϑ is performed over L generations: In each generation the $(\mu/\mu_I, \lambda)$-CMSA-ES generates λ offspring with individual mutation strengths σ_l, see lines 4 to 10. The mutation strength σ_l can be interpreted as an individual scaling factor that is self-adaptively evolved using the learning parameter $\tau_\sigma = \frac{1}{\sqrt{2N}}$ (N – search space dimension). The mutation vector $\mathbf{z}_l$ of each offspring depends on the covariance matrix $\mathbf{C}$ which corresponds to the distribution of previously generated successful candidate solutions. The update rule can be found in line 30 where $\tau_c = 1 + \frac{N(N+1)}{2\mu}$ is used. After creation of the offspring, the objective function (fitness) values are calculated. Having completed the offspring population, the algorithm selects those μ of the λ offspring with the best (noisy) fitness values $\tilde{f}_{m;\lambda}$, $m = 1, \ldots, \lambda$. Notice, $m;\lambda$ denotes the mth best out of λ individuals. Accordingly, the notation $\langle . \rangle$ refers to the construction of the centroid of the respective values corresponding to the μ best offspring solutions. For example, the centroid of the mutation strengths is $\langle \sigma \rangle = \frac{1}{\mu} \sum_{m=1}^{\mu} \sigma_{m;\lambda}$. Subsequently, the pcCMSA-ES examines the list $\mathcal{F}$ using the linear regression approach. The hypothesis test (6) is implemented within the program `detection`$(\mathcal{F}_{int}, \alpha)$, line 19. Analyzing the fitness interval $\mathcal{F}_{int}$, it returns the decision variable $td = 1$ if (6) is fulfilled, else $td = 0$. The parameter α refers to the significance level of the hypothesis test. As long as a negative trend is detected the algorithm acts like the original CMSA-ES. Indication of a non-negative trend ($td = 0$) leads to an increase of the population size μ by multiplication with the factor $c_\mu > 1$, line 21, keeping the truncation ratio $\vartheta = \mu/\lambda$ constant by line 4. In order to prevent the next hypothesis test from being biased by old fitness values, the detection procedure is interrupted for L generations (line 26). Additionally the covariance matrix adaptation in line 30 is turned off, once the algorithm has encountered significant noise impact. For this purpose the parameter $adjC$ is set to zero in line 22. Stalling the covariance matrix update is necessary to avoid a random matrix process resulting in a rise of the condition number of $\mathbf{C}$ without gaining any useful information from the noisy environment.

In the case the hypothesis test returned $td = 1$, i.e. (6) is fulfilled, there is a significant convergence trend. In such a situation one can try to minimize the efforts and reduce the population size in line 24. Such a reduction can make sense in the distance dependent noise case (2a) where there is a minimal population size above which the ES

Algorithm 1 pcCMSA-ES

1: Initialization: $g \leftarrow 0$; $wait \leftarrow 0$; $\langle \sigma \rangle \leftarrow \sigma^{(init)}$; $\langle \boldsymbol{y} \rangle \leftarrow \boldsymbol{y}^{(init)}$;
2: $\quad \mu \leftarrow \mu^{(init)}$; $\mu_{\min} \leftarrow \mu^{(init)}$; $\boldsymbol{C} \leftarrow \boldsymbol{I}$; $adjC \leftarrow 1$
3: **repeat**
4: $\quad \lambda \leftarrow \lfloor \mu/\vartheta \rfloor$
5: $\quad$ **for** $l \leftarrow 1$ to λ **do**
6: $\quad\quad \sigma_l \leftarrow \langle \sigma \rangle \mathrm{e}^{\tau_\sigma \mathcal{N}(0,1)}$
7: $\quad\quad \boldsymbol{s}_l \leftarrow \sqrt{\boldsymbol{C}} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$
8: $\quad\quad \boldsymbol{z}_l \leftarrow \sigma_l \boldsymbol{s}_l$
9: $\quad\quad \boldsymbol{y}_l \leftarrow \langle \boldsymbol{y} \rangle + \boldsymbol{z}_l$
10: $\quad\quad \tilde{f}_l \leftarrow \tilde{f}(\boldsymbol{y}_l)$
11: $\quad$ **end for**
12: $\quad g \leftarrow g + 1$
13: $\quad \langle \boldsymbol{z} \rangle \leftarrow \sum_{m=1}^{\mu} \boldsymbol{z}_{m;\lambda}$
14: $\quad \langle \sigma \rangle \leftarrow \sum_{m=1}^{\mu} \sigma_{m;\lambda}$
15: $\quad \langle \boldsymbol{y} \rangle \leftarrow \langle \boldsymbol{y} \rangle + \langle \boldsymbol{z} \rangle$
16: $\quad$ **add** $\tilde{f}(\langle \boldsymbol{y} \rangle)$ TO $\mathcal{F}$
17: $\quad$ **if** $g > L \wedge wait = 0$ **then**
18: $\quad\quad \mathcal{F}_{int} \leftarrow \mathcal{F}(g - L : g)$
19: $\quad\quad td \leftarrow \mathtt{detection}(\mathcal{F}_{int}, \alpha)$
20: $\quad\quad$ **if** $td = 0$ **then**
21: $\quad\quad\quad \mu \leftarrow \mu c_\mu$
22: $\quad\quad\quad adjC \leftarrow 0$
23: $\quad\quad$ **else**
24: $\quad\quad\quad \mu \leftarrow \max\left(\mu_{\min}, \lfloor \mu/b_\mu \rfloor\right)$
25: $\quad\quad$ **end if**
26: $\quad\quad wait \leftarrow L$
27: $\quad$ **else if** $wait > 0$ **then**
28: $\quad\quad wait \leftarrow wait - 1$
29: $\quad$ **end if**
30: $\quad \boldsymbol{C} \leftarrow \left(1 - \frac{1}{\tau_c}\right)^{adjC} \boldsymbol{C} + \frac{adjC}{\tau_c} \langle \boldsymbol{s}\boldsymbol{s}^\top \rangle$
31: **until** termination condition
32: **return** $\langle \boldsymbol{y} \rangle$

converges without further population size increase. That is, the pcCMSA-ES increases first the population size aggressively and after reaching convergence, the population size is slowly decreased to its nearly optimal value. That is, the reduction factor b_μ should be related to that of c_μ, e.g. $b_\mu = \sqrt[k]{c_\mu}$ ($k = 2$, or 3), or can be chosen independently, but should fulfill $b_\mu < c_\mu$.

Regarding fitness environments where the ES has to deal temporarily with noisy regions, it might be beneficial to turn the covariance matrix adaptation on again once the ES has left the noisy region. That is, if a significant negative trend is present again the parameter $adjC$ should be reset to one in order to gain additional information about advantageous search directions. This can easily be obtained by inserting $adjC \leftarrow 1$ after line 24. However, as for the noisy fitness environments considered here, this adjustment is not able to provide significant improvements in terms of the ES's progress and therefore has not been implemented. It remains to be investigated in further studies.

4 Experimental Investigations and Discussion

The behavior of the proposed pcCMSA-ES algorithm is investigated on the ellipsoid model

$$f(\mathbf{y}) = \sum_{i=1}^{N} q_i y_i^2 \tag{7}$$

with noise types (2b) and (2a). Especially, the cases $q_i = 1$ (sphere model) and $q_i = i, i^2$ have been considered. In the simulations the pcCMSA-ES is initialized with standard parameter settings and $\sigma^{(init)} = 1$ at $\mathbf{y}^{(init)} = \mathbf{1}$ in search space dimension $N = 30$. The initial population sizes are set to $\mu = 3$ and $\lambda = 9$ resulting in a truncation ratio $\vartheta = \frac{\mu}{\lambda} = \frac{1}{3}$ during the runs. The population size factors are $c_\mu = 2$ and $b_\mu = \sqrt{c_\mu}$. The significance level of the hypothesis test in line 19 of Alg. 1 is $\alpha = 0.05$. The length L of the $\tilde{f}$-data collection phases must be chosen long enough to ensure a sufficient f improvement. As shown in [10], the effort to get an expected relative f improvement is proportional to the quotient of the trace of the Hessian of f and its minimal eigenvalue. Hence, for the sphere the effort is proportional to N and for the $q_i = i^2$ ellipsoid proportional to $\Sigma q := \sum_{i=1}^{N} q_i$. In the experiments $L = 5N$ and $L = \Sigma q$ are used.

Figure 1 shows the pcCMSA-ES dynamics for the (2b) case of constant $\sigma_\epsilon = 1$ noise. Considering the simple regret curves (blue), after a transient phase one observes that the ES on average continuously approaches the optimizer at a linear order in the log-log-plot. That means that $\mathrm{SR}(n) \propto n^a$ with $a < 0$. The parallelly descending dashed (magenta) lines $h(n) \propto n^{-1}$ indicate that the pcCMSA-ES actually realizes an $a \approx -1$. Fitting linear curves (solid magenta) to those $\mathrm{SR}(n)$ graphs, using the technique described by Eqs. (3)–(5), one can calculate the confidence intervals for a given confidence level, e.g. 95%, which is displayed in Fig. 1. The observed $a \approx -1$ is remarkable since it apparently seems to violate a theorem by Astete-Morales, Cauwet, and Teytaud [7] that states that "Simple ES" can only reach an $a > -\frac{1}{2}$. The authors even supported their theorem with experiments regarding a tailored (1 + 1)-ES with resampling that came close to $-\frac{1}{2}$ and the UH-CMA-ES [5] that produced only a-values in the range of -0.1 to -0.3. Having a look at the assumptions made to prove the theorem, one finds the reason in the definition of "Simple ES". It contains a common assumption regarding the operation of ES – the scale invariance of the mutations. Roughly speaking, the expected value of the mutation strength should scale with the distance to the optimizer. That is, if one gets closer to the optimizer, the mutation strength should shrink. Looking at the (green) $\langle\sigma\rangle$ dynamics in Fig. 1 one sees that this assumption does not hold for the pcCMSA-ES. Remarkably, $\langle\sigma\rangle$ reaches a constant steady state value. Since theorems cannot be wrong, unlike the (1 + 1)-ES and the UH-CMA-ES, the pcCMSA-ES is *not* a "Simple ES".

While the pcCMSA-ES approaches a fixed mutation strength, on average it approaches the optimizer continuously as can be seen in Fig. 1 where the dynamics of the weighted residual distance R_q to the optimizer is displayed (red curves). This distance measure is defined as $R_q(\mathbf{y}) := \sqrt{\sum_{i=1}^{N} q_i^2 y_i^2}$. According to formula (22) in [3] the steady state expected value of $R_q(\mathbf{y})$ can be estimated for fixed population sizes

$$R_q^{\mathrm{ss}} = \sqrt{\frac{\sigma_\epsilon \Sigma q}{4\mu c_{\mu/\mu,\lambda}}}, \tag{8}$$

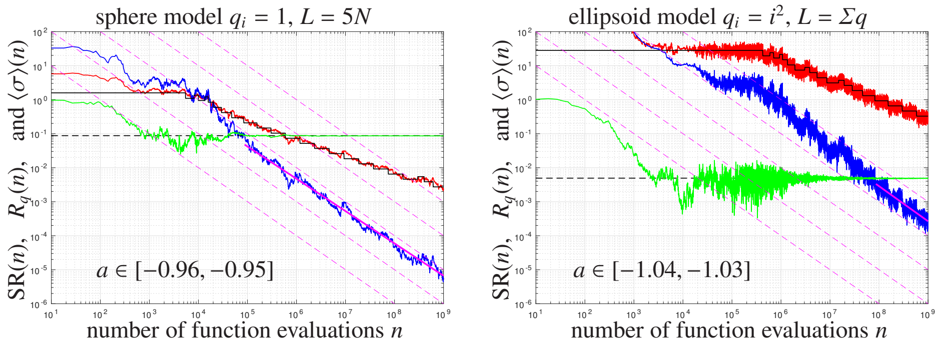

Fig. 1. The dynamical behavior of the pcCMSA-ES subject to additive fitness noise of strength $\sigma_\epsilon = 1$. Considering the sphere model as well as the ellipsoid model $q_i = i^2$ and search space dimensionality $N = 30$, four dynamics are plotted against the number of function evaluations n: the noise-free fitness of the parental centroid $\mathrm{SR}(n) = f(\langle \mathbf{y} \rangle)$ (blue), the corresponding weighted residual distance $R_q(n) = R_q(\langle \mathbf{y} \rangle)$ (red), and the mutation strength $\langle \sigma \rangle$ (green). The solid black step function predicts the residual steady state distance according to Eq. (8). In both cases, it is steadily reduced with each μ elevation.

where $c_{\mu/\mu,\lambda}$ is the well-known progress coefficient [11]. This distance is reached by the CMSA-ES after a sufficiently long generation period (keeping μ and λ constant). Since the pcCMSA-ES changes the population size successively, the theoretical estimate (8) can be used to check whether the population size dynamics of the pcCMSA-ES works satisfactorily. The R_q dynamics follows closely the prediction of (8), which are displayed as (black) staircase curves.

As a second example, the case of distance dependent noise is considered in Fig. 2. The noise variance vanishes when approaching the optimizer. According to the progress rate theory for the noisy ellipsoid [12], one can derive an evolution condition

$$4\mu^2 c_{\mu/\mu,\lambda}^2 > \sigma^{*2} + \sigma_\epsilon^{*2} \tag{9}$$

that states that given upper values of normalized normalized noise and mutation strengths there is a parental population size μ (μ/λ = const.) above which the ES converges to the optimizer. Here the normalized quantities are defined as $\sigma^* := \sigma \Sigma q / R_q$ and $\sigma_\epsilon^* := \sigma_\epsilon \Sigma q / (2R_q^2)$. Figure 2 shows the dynamics of the pcCMSA-ES on sphere and ellipsoid ($q_i = i^2$) model with normalized noise strengths $\sigma_\epsilon^* = 10$ and $\sigma_\epsilon^* = 4$, respectively. Taking a look at the solid blue lines representing the simple regret (being the noise-free fitness dynamics $f(\langle \mathbf{y} \rangle)$), one observes initially an *increase* of the parental simple regret. That is, the pcCMSA-ES departs from the optimizer. This is due to the choice of the initial population size of $\mu = 3$, $\lambda = 9$ being too small. However, after the first L generations, the first hypothesis test indicates divergence and the population size μ is increased by a factor $c_\mu = 2$. This increase repeats two or three times, as can be seen considering the (black) staircase curves displaying λ in Fig. 2, until a population size has been reached where the hypothesis test in line 19 of Alg. 1 returns 1 indicating convergence, the SR-curves start to descend. This behavior is also reflected by the dynamics of the residual distance to the optimizer $R_q(\langle \mathbf{y} \rangle)$ (red). This attests that

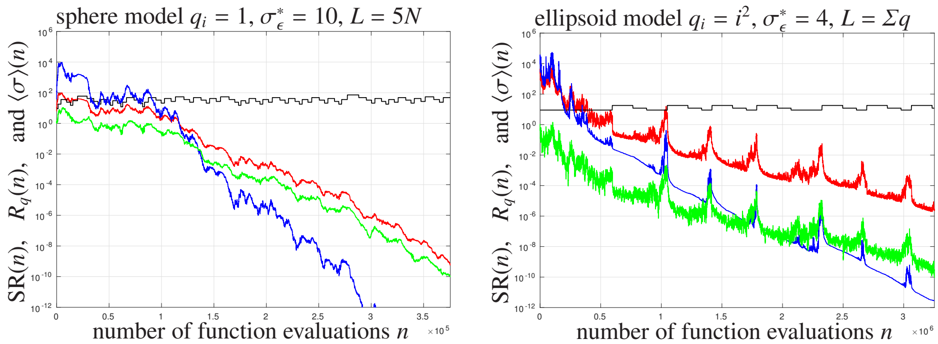

Fig. 2. The dynamical behavior of the pcCMSA-ES subject to distance dependent noise of normalized noise strength σ^*_ϵ. Considering the sphere model as well as the ellipsoid model $q_i = i^2$ with search space dimensionality $N = 30$, four dynamics are plotted against the number of function evaluations n: the simple regret of the parental centroid $\langle \boldsymbol{y} \rangle$ (blue), the corresponding residual distance $R_q(\langle \boldsymbol{y} \rangle)$ (red), and the mutation strength $\langle \sigma \rangle$ (green). The solid black staircase presents the offspring population size $\lambda = \lfloor \mu/\vartheta \rfloor$. According to Eq. (9), it will be increased up to a value where the strategy is able to establish continuous progress towards the optimizer. Afterwards the population size fluctuates around that specific value.

the pcCMSA-ES is able to adapt an appropriate population size needed to comply with Eq. (9) rather than simply increasing it arbitrarily. In contrast to the previous case of additive noise the mutation strength dynamics in Fig. 2 indicate a successive reduction of the noise strength σ. This is due to the decreasing influence of the distance dependent noise as the ES approaches the optimizer. In such cases the behavior of a "Simple ES" is desirable. The pcCMSA-ES behaves as such and demonstrates its ability to exhibit a linear convergence order similar to the non-noisy case. However, it has to be pointed out that the current population size reduction rule can result in interrupted convergence behavior in cases of very strong distance dependent noise. This can be inferred from the peaks in the right graph of Fig. 2. An attempt to address this disruption would be shortening both the test interval length L as well as the waiting time *wait* of the algorithm after each population size reduction and enlarging them again after a population size escalation, respectively. Also switching off the population size reduction might be a reasonable approach. Eventually, the population size control configuration under severe fitness proportional noise should be examined more closely in future investigations.

5 Summary and Outlook

This paper presented an EA for the treatment of noisy optimization problems that is based on the CMSA-ES. Within its concept a mechanism for identification of noise-related stagnations or divergence behavior is integrated. Consequently, having identified noise related behavior the algorithm increases the size of the parental as well as the offspring population. This way it improves the likelihood to approach closer residual distances to the optimizer. Significant noise disturbances become noticeable by the absence of a clearly negative trend (minimization considered) within the noisy fitness dynamics. The slope of the respective trend can be deduced from the corresponding linear regression line. The estimated trend is used in a hypothesis test to decide whether there is convergence to the optimizer. If no further significant noise influences are discovered in subsequent tests the population size is again gradually reduced to avoid unnecessary function evaluations. This way the algorithm is capable to adapt the appropriate populations size. Accordingly, the adjusted CMSA-ES is denoted population control covariance matrix self-adaptation evolution strategy – pcCMSA-ES.

As a proof of concept, the pcCMSA-ES was tested on the noisy ellipsoid model considering two noise models, which obey different characteristics. The additive fitness noise case with constant noise strength σ_ϵ requires a permanent increase of the population size. On the other hand, the distance dependent noise case (which is equivalent to fitness proportionate noise in the case of the sphere model) requires only a limited population size increase. A well-crafted EA should be able to handle both cases (and of course, non-noisy optimization problems as well).

The empirical investigation of the strong noise case $\sigma_\epsilon = \text{const.}$ revealed a remarkable behavior of the pcCMSA-ES. The dynamics by which this ES approaches the optimizer seems to be already the fastest one can expect from a direct search algorithm on quadratic functions. The simple regret obeys an n^a dynamics with $a \approx -1$. This is remarkable since "Simple ES" should only allow for an $a \geq -1/2$ no matter how the noise is handled. The reason for this behavior is that unlike "Simple ES" the pcCMSA-ES does not scale the mutation strength σ in proportion to the distance to the optimizer in case of strong noise. This is different to other ESs such as $(1 + 1)$ or UH-CMA. However, if there is no strong noise, pcCMSA-ES behaves like a "Simple ES".

The pcCMSA-ES requires the fixing of additional exogenous strategy parameters. Particularly, the length L of the interval of observed fitness values that are considered in a single test decision has to be examined more closely. L should be large enough to ensure a sufficient evolution (convergence) of the fitness values. From the progress rate theory it is known that the number of generations needed for a certain fitness improvement scale with the quotient of the trace of the Hessian of f and its smallest eigenvalue. Therefore, L should be chosen proportional to N (search space dimensionality) in the sphere model case and to $\frac{N}{6}(N + 1)(2N + 1)$ in the case of the ellipsoid model $q_i = i^2$. However, in the black-box scenario the Hessian is not known. However, as long as the initial noise influence is small, the pcCMSA-ES transforms the optimization problem gradually into a local sphere model. In such cases, the $L \propto N$ choice should suffice. If, however, the noise is already strong in the initial phase, there is no definitive choice and the user has to make a guess regarding the trace vs. minimum eigenvalue ratio. Choosing L too large has a negative influence on the efficiency of the ES. It effects the lead time of the algorithm needed to establish an initial interval of fitness observations $\mathcal{F}_{int}$ as well as the waiting time *wait*. The parameter *wait* governs the length of the waiting period after a single population adjustment. After a transient phase of *wait* generations the algorithm starts again with the analysis of the fitness dynamics. It is not evident whether the parameter *wait* should depend on the length L of the fitness interval. The waiting time is essential to prevent wrong test decisions based on fitness dynamics resulting from different population specifications. A beneficial parameter setting has to be determined in future empirical investigations. There are also open questions regarding a profound choice of the population size change parameters c_μ and b_μ and the significance level $\alpha = 0.05$ used. These question should be also tackled by extended empirical investigations considering different test functions and noise scenarios.

Regarding theory, the analysis of certain aspects of the pcCMSA-ES seems to be possible using and extending the results presented in [12]. For example, the observed steady state σ in the strong noise case should be deducible from the self-adaptation response theory. Deriving the remarkable empirically observed $\mathrm{SR}(n) \propto n^{-1}$ law is clearly another task for future research.

Acknowledgements

This work was supported by the Austrian Science Fund FWF under grant P22649-N23 and by the Austrian funding program COMET (COMpetence centers for Excellent Technologies) in the K-Project Advanced Engineering Design Automation (AEDA).

References

[1] D. V. Arnold. *Noisy Optimization with Evolution Strategies*. Kluwer, 2002.

[2] Y. Jin and J. Branke. Evolutionary optimization in uncertain environments - a survey. *IEEE Trans. Evolutionary Computation*, 9(3):303–317, 2005.

[3] H.-G. Beyer, D. V. Arnold, and S. Meyer-Nieberg. A New Approach for Predicting the Final Outcome of Evolution Strategy Optimization Under Noise. *Genetic Programming and Evolvable Machines*, 6(1):7–24, 2005.

[4] H.-G. Beyer and B. Sendhoff. Evolution Strategies for Robust Optimization. In *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on*, pages 1346–1353, 2006.

[5] N. Hansen, A. S. P. Niederberger, L. Guzzella, and P. Koumoutsakos. A Method for Handling Uncertainty in Evolutionary Optimization With an Application to Feedback Control of Combustion. *IEEE Trans. Evolutionary Computation*, 13(1):180–197, 2009.

[6] H.-G. Beyer and B. Sendhoff. Covariance Matrix Adaptation Revisited - The CMSA Evolution Strategy. In Günter Rudolph, Thomas Jansen, Nicola Beume, Simon Lucas, and Carlo Poloni, editors, *Parallel Problem Solving from Nature - PPSN X*, volume 5199 of *Lecture Notes in Computer Science*, pages 123–132. Springer Berlin Heidelberg, 2008.

[7] S. Astete-Morales, M.-L. Cauwet, and O. Teytaud. Evolution strategies with additive noise: A convergence rate lower bound. In *Proceedings of the 2015 ACM Conference on Foundations of Genetic Algorithms XIII*, FOGA '15, pages 76–84, New York, NY, USA, 2015. ACM.

[8] O. Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *COLT 2013 - The 26th Annual Conference on Learning Theory, NJ, USA*, pages 3–24, 2013.

[9] J. F. Kenney. *Mathematics of Statistics*. Number 1-2. Literary Licensing, LLC, 2013.

[10] H.-G. Beyer and A. Melkozerov. The Dynamics of Self-Adaptive Multi-Recombinant Evolution Strategies on the General Ellipsoid Model. *IEEE Transactions on Evolutionary Computation*, 18(5):764–778, 2014.

[11] H.-G. Beyer. *The Theory of Evolution Strategies*. Natural Computing Series. Springer, Heidelberg, 2001.

[12] A. Melkozerov and H.-G. Beyer. Towards an Analysis of Self-Adaptive Evolution Strategies on the Noisy Ellipsoid Model: Progress Rate and Self-Adaptation Response. In *GECCO 2015, Madrid, Spain*, pages 297–304. ACM, 2015.