



Variance Reduction for Better Sampling in Continuous Domains

Laurent Meunier^{1,2(✉)}, Carola Doerr³, Jeremy Rapin¹, and Olivier Teytaud¹

¹ Facebook Artificial Intelligence Research (FAIR), Paris, France

laurentmeunier@fb.com

² PSL, Université Paris-Dauphine, Miles Team, Paris, France

³ Sorbonne Université, CNRS, LIP6, Paris, France

Abstract. Design of experiments, random search, initialization of population-based methods, or sampling inside an epoch of an evolutionary algorithm uses a sample drawn according to some probability distribution for approximating the location of an optimum. Recent papers have shown that the optimal *search* distribution, used for the sampling, might be more peaked around the center of the distribution than the *prior* distribution modelling our uncertainty about the location of the optimum. We confirm this statement, provide explicit values for this reshaping of the search distribution depending on the population size λ and the dimension d , and validate our results experimentally.

1 Introduction

We consider the setting in which one aims to locate an optimal solution $x^* \in \mathbb{R}^d$ for a given black-box problem $f : \mathbb{R}^d \rightarrow \mathbb{R}$ through a parallel evaluation of λ solution candidates. A simple, yet effective strategy for this *one-shot optimization* setting is to choose the λ candidates from a normal distribution $\mathcal{N}(\mu, \sigma^2)$, typically centered around an *a priori* estimate μ of the optimum and using a variance σ^2 that is calibrated according to the uncertainty with respect to the optimum. Random independent sampling is – despite its simplicity – still a very commonly used and performing good technique in one-shot optimization settings. There also exist more sophisticated sampling strategies like Latin Hypercube Sampling (LHS [19]), or quasi-random constructions such as Sobol, Halton, Hammersley sequences [7, 18] – see [2, 6] for examples. However, no general superiority of these strategies over random sampling can be observed when the benchmark set is sufficiently diverse [4]. It is therefore not surprising that in several one-shot settings – for example, the design of experiments [1, 13, 19, 21] or the initialization (and sometimes also further iterations) of evolution strategies – the solution candidates are frequently sampled from random independent distributions (though sometimes improved by mirrored sampling [27]). A surprising finding was recently communicated in [6], where the authors consider the setting in which the optimum x^* is known to be distributed according to

Arxiv version [20] of the present document includes bigger plots and the appendices.

© Springer Nature Switzerland AG 2020

T. Bäck et al. (Eds.): PPSN 2020, LNCS 12269, pp. 154–168, 2020.

https://doi.org/10.1007/978-3-030-58112-1_11




Fig. 1. Average regret, normalized by d , on the sphere function for various dimensions and budgets in terms of rescaled standard deviation. Each mean has been estimated from 100,000 samples. Table on the right: Average regret for $\sigma^* = \sqrt{\log(\lambda)/d}$ and $\sigma = 1$.

a standard normal distribution $\mathcal{N}(0, I_d)$, and the goal is to minimize the distance of the best of the λ samples to this optimum. In the context of evolution strategies, one would formulate this problem as minimizing the sphere function with a normally distributed optimum. Intuitively, one might guess that sampling the λ candidates from the same prior distribution, $\mathcal{N}(0, I_d)$, should be optimal. This intuition, however, was disproved in [6], where it is shown that – unless the sample size λ grows exponentially fast in the dimension d – the median quality of sampling from $\mathcal{N}(0, I_d)$ is worse than that of sampling a single point, namely the center point 0. A similar observation was previously made in [22], without mathematically proven guarantees.

Our Theoretical Result. It was left open in [6] how to optimally scale the variance σ^2 when sampling the λ solution candidates from a normal distribution $\mathcal{N}(0, \sigma^2 I_d)$. While the result from [6] suggests to use $\sigma = 0$, we show in this work that a more effective strategy exists. More precisely, we show that setting $\sigma^2 = \min\{1, \Theta(\log(\lambda)/d)\}$ is asymptotically optimal, as long as λ is sub-exponential, but growing in d . Our variance scaling factor reduces the median approximation error by a $1 - \varepsilon$ factor, with $\varepsilon = \Theta(\log(\lambda)/d)$. We also prove that no constant variance nor any other variance scaling as $\omega(\log(\lambda)/d)$ can achieve such an approximation error. Note that several optimization algorithms operate with rescaled sampling. Our theoretical results therefore set the mathematical foundation for empirical rules of thumb such as, for example, used in e.g. [6, 8–10, 17, 22, 28].

Our Empirical Results. We complement our theoretical analyses by an empirical investigation of the rescaled sampling strategy. Experiments on the sphere

function confirm the results. We also show that our scaling factor for the variance yields excellent performance on two other benchmark problems, the Cigar and the Rastrigin function. Finally, we demonstrate that these improvements are not restricted to the one-shot setting by applying them to the initialization of iterative optimization strategies. More precisely, we show a positive impact on the initialization of Bayesian optimization algorithms [15] and on differential evolution [25].

Related Work. While the most relevant works for our study have been mentioned above, we briefly note that a similar surprising effect as observed here is the “Stein phenomenon” [14, 24]. Although an intuitive way to estimate the mean of a standard gaussian distribution is to compute the empirical mean, Stein showed that this strategy is sub-optimal w.r.t. mean squared error and that the empirical mean needs to be rescaled by some factor to be optimal.

2 Problem Statement and Related Work

The context of our theoretical analysis is *one-shot optimization*. In one-shot optimization, we are allowed to select λ points $x_1, \dots, x_\lambda \in \mathbb{R}^d$. The quality $f(x_i)$ of these points is evaluated, and we measure the performance of our samples in terms of simple regret [5] $\min_{i=1, \dots, \lambda} f(x_i) - \inf_{x \in \mathbb{R}^d} f(x)$.¹ That is, we aim to minimize the distance – measured in *quality space* – of the best of our points to the optimum. This formulation, however, also covers the case in which we aim to minimize the distance to the optimum in the *search space*: we simply take as f the root of the sphere function $f_{x^*} : \mathbb{R}^d \rightarrow \mathbb{R}, x \mapsto \|x - x^*\|^2$, where here and in the following $\|\cdot\|$ denotes the Euclidean norm.

Rescaled Random Sampling for Randomly Placed Optimum. In the setting studied in Sect. 3 we assume that the optimum x^* is sampled from the standard multivariate Gaussian distribution $\mathcal{N}(0, I_d)$, and that we aim to minimize the regret $\min_{i=1, \dots, \lambda} \|x_i - x^*\|^2$ through i.i.d. samples $x_i \sim \mathcal{N}(0, \sigma^2 I_d)$. That is, in contrast to the classical *design of experiments* (DoE) setting, we are only allowed to choose the scaling factor σ , whereas in DoE more sophisticated (often quasi-random and space-filling designs – which are typically not i.i.d. samples) are admissible. Intuitively, one might be tempted to guess that $\sigma = 1$ should be a good choice, as in this case the λ points are chosen from the same distribution as the optimum x^* . This intuition, however, was refuted in [6, Theorem 1], where it was shown that the middle point sampling strategy, which uses $\sigma = 0$ (i.e., all λ points collapse to $(0, \dots, 0)$) yields smaller regret than sampling from $\mathcal{N}(0, I_d)$ unless λ grows exponentially in d . More precisely, it is shown in [6] that, for this regime of λ and d , the median of $\|x^*\|^2$ is smaller than the median of $\|x_i - x^*\|^2$

¹ This requires knowledge of $\inf_x f(x)$, which may not be available in real-world applications. In this case, without loss of generality (this is just for the sake of plotting regret values), the infimum can be replaced by an empirical minimum. In all applications considered in this work the value of $\inf_x f(x)$ is known.




Fig. 2. Comparison of methods: without rescaling ($\sigma = 1$), middle point sampling ($\sigma = 0$), and our rescaling method ($\sigma = \sqrt{\frac{\log \lambda}{d}}$). Each mean has been estimated from 10^5 samples. (On left) Average regret, normalized by d , on the sphere function for diverse population sizes λ at fixed dimension $d = 20$. The gain of rescaling decreases as λ increases. (On right) Distribution of the regret for the strategies on the $50d$ -sphere function for $\lambda = 1000$.

for i.i.d. $x_i \in \mathcal{N}(0, I_d)$. This shows that sampling a single point can be better than sampling λ points with the wrong scaling factor, unless the budget λ is very large.

Our goal is to improve upon the middle point strategy, by deriving a scaling factor σ such that the λ i.i.d. samples yield smaller regret with a decent probability. More precisely, we aim at identifying σ such that

$$\mathbb{P} \left[\min_{1 \leq i \leq \lambda} \|x_i - x^*\|^2 \leq (1 - \varepsilon) \|x^*\|^2 \right] \geq \delta, \quad (1)$$

for some $\delta \geq 1/2$ and $\varepsilon > 0$ as large as possible. Here, in line with [6], we have switched to regret, for convenience of notation. [6] proposed, without proof, such a scaling factor: our proposal is dramatically better in some regimes.

3 Theoretical Results

We derive sufficient and necessary conditions on the scaling factor σ such that Eq. (1) can be satisfied. More precisely, we prove that Eq. (1) holds with approximation gain $\varepsilon \approx \log(\lambda)/d$ when the variance σ^2 is chosen proportionally to $\log \lambda/d$ (and λ does not grow too rapidly in d). We then show that Eq. (1) cannot be satisfied for $\sigma^2 = \omega(\log(\lambda)/d)$. Moreover, we prove that $\varepsilon = O(\log(\lambda)/d)$, which, together with the first result, shows that our scaling factor is asymptotically optimal. The precise statements are summarized in Theorems 1, 2, and 3, respectively. Proof sketches are available in Sect. 3. Proofs are left in the full version available on the ArXiv version [20].

Theorem 1 (*Sufficient condition on rescaling*). *Let $\delta \in [\frac{1}{2}, 1)$. Let $\lambda = \lambda_d$, satisfying*

$$\lambda_d \rightarrow \infty \text{ as } d \rightarrow \infty \text{ and } \log(\lambda_d) \in o(d). \quad (2)$$

Then there exist two positive constants c_1, c_2 , and d_0 , such that for all $d \geq d_0$ it holds that

$$\mathbb{P} \left[\min_{i=1, \dots, \lambda} \|x^* - x_i\|^2 \leq (1 - \varepsilon) \|x^*\|^2 \right] \geq \delta \tag{3}$$

when x^* is sampled from the standard Gaussian distribution $\mathcal{N}(0, I_d)$, x_1, \dots, x_λ are independently sampled from $\mathcal{N}(0, \sigma^2 I_d)$ with $\sigma^2 = \sigma_d^2 = c_2 \log(\lambda)/d$ and $\varepsilon = \varepsilon_d = c_1 \log(\lambda)/d$.




Fig. 3. Comparison of various one-shot optimization methods from the point of view of the simple regret. Reading guide in Sect. 4.2. Results are averaged over objective functions Cigar, Rastrigin, Sphere in dimension 20, 200, 2000, and budget 30, 100, 3000, 10000, 30000, 100000. **MetaTuneRecentering** performs best overall. Only the 30 best performing methods are displayed as columns, and the 6 best as rows. Red means superior performance of row vs col. Rows and cols ranked by performance. (Color figure online)

Theorem 1 shows that i.i.d. Gaussian sampling can outperform the middle point strategy derived in [6] (i.e., the strategy using $\sigma^2 = 0$) if the scaling factor σ is chosen appropriately. Our next theorem summarizes our findings for the conditions that are *necessary* for the scaling factor σ^2 to outperform this middle point strategy. This result, in particular, illustrates why neither the natural choice $\sigma = 1$, nor any other constant scaling factor can be optimal.

Theorem 2 (Necessary condition on rescaling). Consider $\lambda = \lambda_d$ satisfying assumptions (2). There exists an absolute constant $C > 0$ such that for all $\delta \in [\frac{1}{2}, 1)$, there exists $d_0 > 0$ such that, for all $d > d_0$ and for all σ the property

$$\exists \varepsilon > 0, \mathbb{P} \left[\min_{i=1, \dots, \lambda} \|x^* - x_i\|^2 \leq (1 - \varepsilon) \|x^*\|^2 \right] \geq \delta \tag{4}$$

for $x^* \sim \mathcal{N}(0, I_d)$ and x_1, \dots, x_λ independently sampled from $\mathcal{N}(0, \sigma^2 I_d)$, implies that $\sigma^2 \leq C \log(\lambda)/d$.

While Theorem 2 induces a necessary condition on the scaling factor σ to improve over the middle point strategy, it does not bound the gain that one can achieve through a proper scaling. Our next theorem shows that the factor derived in Theorem 1 is asymptotically optimal.

Theorem 3 (*Upper bound for the approximation factor*). *Consider $\lambda = \lambda_d$ satisfying assumptions (2). There exists an absolute constant $C' > 0$ such that for all $\delta \in [\frac{1}{2}, 1)$, there exists $d_0 > 0$ such that, for all $d > d_0$ and for all $\varepsilon, \sigma > 0$, it holds that if $\mathbb{P}[\min_{i=1, \dots, \lambda} \|x^* - x_i\|^2 \leq (1 - \varepsilon) \|x^*\|^2] \geq \delta$ for $x^* \sim \mathcal{N}(0, I_d)$ and x_1, \dots, x_λ independently sampled from $\mathcal{N}(0, \sigma^2 I_d)$, then $\varepsilon \leq C' \log(\lambda)/d$.*

Proof Sketches. We first notice that as x^* is sampled from a standard normal distribution $\mathcal{N}(0, I_d)$, its norm satisfies $\|x^*\|^2 = d + o(d)$ as $d \rightarrow \infty$. We then use that, conditionally to x^* , it holds that

$$\mathbb{P}[\min_{i \in [\lambda]} \|x^* - x_i\|^2 \leq (1 - \varepsilon) \|x^*\|^2 | x^*] = 1 - (1 - \mathbb{P}[\|x - x^*\|^2 \leq (1 - \varepsilon) \|x^*\|^2 | x^*])^\lambda$$

We therefore investigate when the condition

$$\mathbb{P}[\|x - x^*\|^2 \leq (1 - \varepsilon) \|x^*\|^2 | x^*] > 1 - (1 - \delta)^{\frac{1}{\lambda}} \quad (5)$$

is satisfied. To this end, we make use of the fact that the squared distance $\|x^*\|^2$ of x^* to the middle point 0 follows the central $\chi^2(d)$ distribution, whereas, for a given point $x^* \in \mathbb{R}^d$, the distribution of the squared distance $\|x - x^*\|^2/\sigma^2$ for $x \sim \mathcal{N}(0, \sigma^2 I_d)$ follows the non-central $\chi^2(d, \mu)$ distribution with non-centrality parameter $\mu := \|x^*\|^2/\sigma^2$. Using the concentration inequalities provided in [29, Theorem 7] for non-central χ^2 distributions, we then derive sufficient and necessary conditions for condition (5) to hold. With this, and using assumptions (2), we are able to derive the results from Theorems 1, 2, and 3.

4 Experimental Performance Comparisons

The theoretical results presented above are in asymptotic terms, and do not specify the constants. We therefore complement our mathematical investigation with an empirical analysis of the rescaling factor. Whereas results for the setting studied in Sect. 3 are presented in Sect. 4.1, we show in Sect. 4.2 that the advantage of our rescaling factor is not limited to minimizing the distance in search space. More precisely, we show that the rescaled sampling achieves good results also in a classical DoE task, in which we aim for minimizing the regret for the Cigar and for the Rastrigin functions. Finally, we investigate in Sect. 4.3 the impact of initializing two common optimization heuristics, Bayesian Optimization (BO) and differential evolution (DE), by a population sampled from the Gaussian distribution $\mathcal{N}(0, \sigma^2 I_d)$ using our rescaling factor $\sigma = \sqrt{\log(\lambda)/d}$.

4.1 Validation of Our Theoretical Results on the Sphere Function

Figure 1 displays the normalized average regret $\frac{1}{d} \mathbb{E} [\min_{i=1, \dots, \lambda} \|x^* - x_i\|^2]$ in terms of $\sigma / \sqrt{\log(\lambda)/d}$ for different dimensions and budgets. We observe that the best parametrization of σ is around $\sqrt{\log(\lambda)/d}$ in all displayed cases. Moreover, we also see that – as expected – the gain of the rescaled sampling over the middle point sampling ($\sigma = 0$) goes to 0 as $d \rightarrow \infty$ (i.e. we get a result closer to the case $\sigma = 0$ as dimension goes to infinity). We also see that, for the regimes plotted in Fig. 1, the advantage of the rescaled variance grows with the budget λ . Figure 2 (on left) displays the average regret (average over multiple samplings and multiple positions of the optimum) as a function of increasing values of λ for the different rescaling methods ($\sigma \in \{0, \sqrt{\log \lambda / d}, 1\}$). We remark, unsurprisingly, that the gain of rescaling is diminishing as $\lambda \rightarrow \infty$. Finally, Fig. 2 (on right) shows the distribution of regrets for the different rescaling methods. The improvement of the expected regret is not at the expense of a higher dispersion of the regret.




Fig. 4. Same experiment as Fig. 3, but separately over each objective function. Results are still averaged over 6 distinct budgets (30, 100, 3000, 10000, 30000, 100000) and 3 distinct dimensionalities (20, 200, 2000). **MetaTuneRecentering** performs well in each case, and is not limited to the sphere function for which it was derived. Variants of LHS are sometimes excellent and sometimes not visible at all (only the 30 best performing methods are shown).

4.2 Comparison with the DoEs Available in Nevergrad

Motivated by the significant improvements presented above, we now investigate whether the advantage of our rescaling factor translates to other optimization tasks. To this end, we first analyze a DoE setting, in which an underlying (and typically not explicitly given) function f is to be minimized through a parallel evaluation of λ solution candidates x_1, \dots, x_λ , and regret is measured in terms of $\min_i f(x_i) - \inf_x f(x)$. In the broader machine learning literature, and in particular in the context of hyper-parameter optimization, this setting is often referred to as *one-shot optimization* [2, 6].

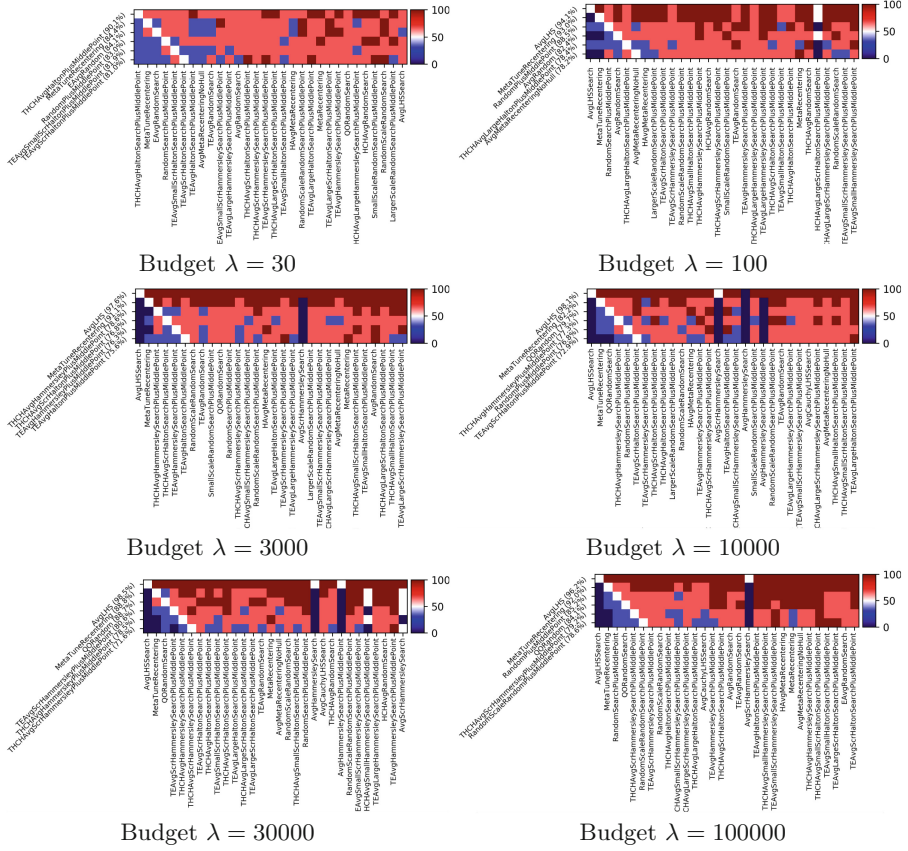


Fig. 5. Methods ranked by performance on the sphere function, per budget. Results averaged over dimension 20, 200, 2000. `MetaTuneRecentering` performs among the best in all cases. LHS is excellent on this very simple setting, namely the sphere function.




Fig. 6. Results on the sphere function, per dimensionality. Results are averaged over 6 values of the budget: 30, 100, 3000, 10000, 30000, 100000. Our method becomes better and better as the dimension increases.

Experimental Setup. All our experiments are implemented and freely available in the Nevergrad platform [23]. Results are presented as shown in Fig. 3. Typically, the six best methods are displayed as rows. The 30 best performing methods

are presented as columns. The order for rows and for columns is the same: algorithms are ranked by their average winning frequency, measured against all other algorithms in the portfolio. The heatmaps show the fraction of runs in which algorithm x (row) outperformed algorithm y (column), averaged over all settings and all replicas (i.e. random repetitions). The settings are typically sweepings over various budgets, dimensions, and objective functions.² For each tested (algorithm, problem) pair, 20 independent runs are performed: a case with N settings is thus based on a total number of $20 \times N$ runs. The number N of distinct problems is at least 6 and often high in the dozens, hence the minimum number of independent runs is at least 120.

Algorithm Portfolio. Several rescaling methods are already available on Nevergrad. A large fraction of these have been implemented by the authors of [6]; in particular:

- The replacement of one sample by the center. These methods are named “midpointX” or “XPlusMiddlePoint”, where X is the original method that has been modified that way.
- The rescaling factor **MetaRecentering** derived in [6]: $\sigma = \frac{1+\log(\lambda)}{4\log(d)}$.
- The quasi-opposite methods suggested in [22], with prefix “QO”: when x is sampled, then another sample $c - rx$ is added, with r uniformly drawn in $[0, 1]$ and c the center of the distribution.

We also include in our comparison a different type of one-shot optimization techniques, independent of the present work, currently available in the platform: they use the information obtained from the sampled points to recommend a point x that is not necessarily one of the λ evaluated ones. These “one-shot+1” strategies have the prefix “Avg”. We keep all these and all other sampling strategies available in Nevergrad for our experiments. We add to this existing Nevergrad portfolio our own rescaling strategy, which uses the scaling factor derived in Sect. 3; i.e., $\sigma = \sqrt{\log(\lambda)/d}$. We refer to this sampling strategy as **MetaTuneRecentering**, defined below. Both scaling factors **MetaRecentering** [6] and **MetaTuneRecentering** (our equations) are applied to quasirandom sampling (more precisely, scrambled Hammersley [1, 13]) rather than random sampling. We provide detailed specifications of these methods and the most important ones below, whereas we skip the dozens of other methods: they are open sourced in Nevergrad [23].

From $[0, 1]^d$ to Gaussian Quasi-random, Random or LHS Sampling: Random sampling, quasi-random sampling, Latin Hypercube Sampling (or others) have a well known definition in $[0, 1]^d$ (for quasi-random, see Halton [12] or Hammersley [13], possibly boosted by scrambling [1]; for LHS, see [19]). To extend to multidimensional Gaussian sampling, we use that if U is a uniform random variable on $[0, 1]$ and Φ the standard Gaussian CDF, then $\Phi^{-1}(U)$ simulates

² Detailed results for individual settings are available at <http://dl.fbaipublicfiles.com/nevergrad/allxps/list.html>.

a $\mathcal{N}(0, 1)$ distribution. We do so on each dimension: this provides a Gaussian quasi-random, random or LHS sampling.

Then, one can rescale the Gaussian quasi-random sampling with the corresponding factor σ for **MetaRecentering** ($\sigma = \frac{1+\log(\lambda)}{4\log(d)}$ [6]) and **MetaTuneRecentering** ($\sigma = \sqrt{\log(\lambda)/d}$): for $i \leq \lambda$ and $j \leq d$, $x_{i,j} = \sigma\phi^{-1}(h_{i,j})$ where $h_{i,j}$ is the j^{th} coordinate of a i^{th} Scrambled-Hammersley point.

Results for the Full DoE Testbed in Nevergrad. Figure 3 displays aggregated results for the Sphere, the Cigar, and the Rastrigin functions, for three different dimensions and six different budgets. We observe that our **MetaTuneRecentering** strategy performs best, with a winning frequency of 80%. It positively compares against all other strategies from the portfolio, with the notable exception of **AvgLHS**, which, in fact, compares favorably against every single other strategy, but with a lower average winning frequency of 73.6%. Note here that **AvgLHS** is one of the “**oneshot+1**” strategies, i.e., it has not only one more sample, but it is also allowed to sample its recommendation adaptively, in contrast to our fully parallel **MetaTuneRecentering** strategy. It performs poorly in some cases (Rastrigin) and does not make sense as an initialization (Sect. 4.3).

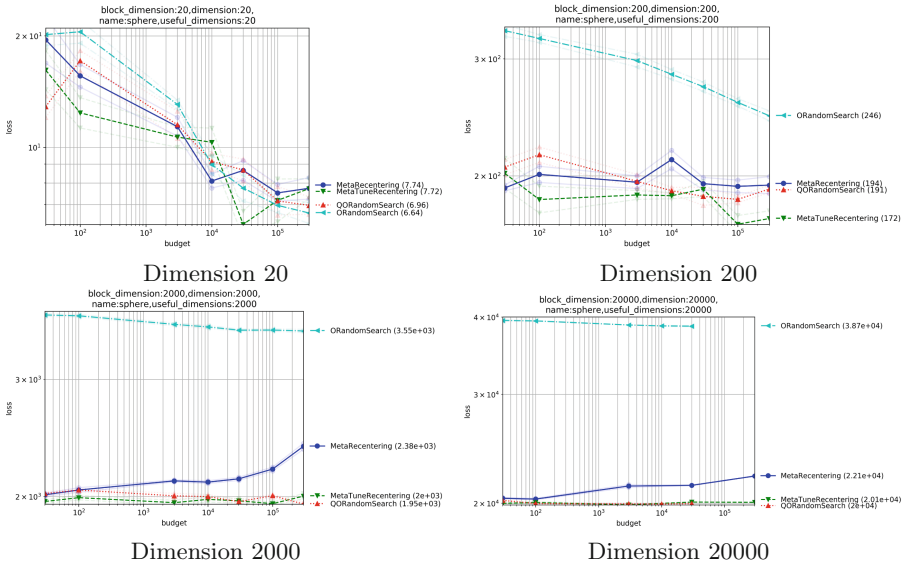


Fig. 7. Same context as Fig. 6, with x -axis = budget and y -axis = average simple regret. We see the failure of **MetaRecentering** in the worsening performance as budget goes to infinity: the budget has an impact on σ which becomes worse, hence worse overall performance. We note that quasi-opposite sampling can perform decently in a wide range of values. Opposite Sampling is not much better than random search in high-dimension. Our **MetaTuneRecentering** shows decent performance: in particular, simple regret decreases as $\lambda \rightarrow \infty$.

Selected DoE Tasks. Figure 4 breaks down the aggregated results from Fig. 3 to the three different functions. We see that **MetaTuneRecentering** scores second on sphere (where AvgLHS is winning), third on Cigar (after AvgLHS and QORandom), and first on Rastrigin. This fine performance is remarkable, given that the portfolio contains quite sophisticated and highly tuned methods. In addition, the AvgLHS methods, sometimes performing better on the sphere, besides using more capabilities than we do (as it is a “oneshot+1” method), had poor results for Rastrigin (not even in the 30 best methods). On sphere, the difference to the third and following strategies is significant (87.3% winning rate against 77.5% for the next runner-up). On Cigar, the differences between the first four strategies are greater than 4% points each, whereas on Rastrigin the average winning frequencies of the first five strategies is comparable, but significantly larger than that of the sixth one (which scores 78.8% against >94.2% for the first five DoEs). Figure 5 zooms into the results for the sphere function, and breaks them further down by available budget λ (note that the results are still averaged over the three tested dimensions). **MetaTuneRecentering** scores second in all six cases. A breakdown of the results for sphere by dimension (and aggregated over the six available budgets) is provided in Fig. 6 and Fig. 7. For dimension 20, we see that **MetaTuneRecentering** ranks third, but, interestingly, the two first methods are “oneshot+1” style (Avg prefix). In dimension 200, **MetaTuneRecentering** ranks second, with considerable advantage over the third-ranked strategy (88.0% vs. 80.8%). Finally, for the largest tested dimension, $d = 2000$, our method ranks first, with an average winning frequency of 90.5%.




Fig. 8. Performance comparison of different strategies to initialize Bayesian Optimization (BO, left) and Differential Evolution (DE, right). A detailed description is given in Sect. 4.3. **MetaTuneRecentering** performs best as an initialization method. In the case of DE, methods different from the traditional DE remain the best on this testcase: when we compare DE with a given initialization and DE initialized with **MetaTuneRecentering**, **MetaTuneRecentering** performs best in almost all cases.

4.3 Application to Iterative Optimization Heuristics

We now move from the one-shot settings considered thus far to *iterative optimization*, and show that our scaling factor can also be beneficial in this context.

More precisely, we analyze the impact of initializing efficient global optimization (EGO [15], a special case of Bayesian optimization) and differential evolution (DE [25]) by a population that is sampled from a distribution that uses our variance scaling scheme. It is well known that a proper initialization can be very critical for the performance of these solvers; see [3, 11, 16, 22, 26] for discussions. Figure 8 summarizes the results of our experiments. As in the previous setups, we compare against existing methods from the Nevergrad platform, to which we have just added our rescaling factor termed **MetaTuneRecentering**. For each initialization scheme, four different initial population sizes are considered: denoting by d the dimension, by w the parallelism (i.e., the number of workers), and by b the total budget that the algorithms can spend on optimizing the given optimization task, the initial population λ is set as $\lambda = \sqrt{b}$ for **Sqrt**, as $\lambda = d$ for **Dim**, $\lambda = w$ for no suffix, and as $\lambda = 30$ when the suffix is **30**. As in Sect. 4.2 we superpose our scaling scheme on top of the quasi-random Scrambled Hammersley sequence suggested in [6], but we also consider random initialization rather than quasi-random (indicated by the suffix “**R**”) and Latin Hypercube Sampling [19] (suffix “**LHS**”). The left chart in Fig. 8 is for the Bayesian optimization case. It aggregates results for 48 settings, which stem from Nevergrad’s “parahdbo4d” suite. It comprises the four benchmark problems Sphere, Cigar, Ellipsoid and Hm. Results are averaged over the total budgets $b \in \{25, 31, 37, 43, 50, 60\}$, dimension $d \in \{20, 2000\}$, and parallelism $w = \max(d, \lfloor b/6 \rfloor)$. We observe that a BO version using our **MetaTuneRecentering** performs best, and that several other variants using this scaling appear among the top-performing configurations. The chart on the right of Fig. 8 summarizes results for Differential Evolution. Since DE can handle larger budgets, we consider here a total number of 100 settings, which correspond to the testcase named “paraalldes” in Nevergrad. In this suite, results are averaged over budgets $b \in \{10, 100, 1000, 10000, 100000\}$, dimensions $d \in \{5, 20, 100, 500, 2500\}$, parallelism $w = \max(d, \lfloor b/6 \rfloor)$, and again the objective functions Sphere, Cigar, Ellipsoid, and Hm. Specialized versions of DE perform best for this testcase, but we see that DE initialized with our **MetaTuneRecentering** strategy ranks fifth (outperformed only by ad hoc variants of DE), with an overall winning frequency that is not much smaller than that of the top-ranked **NoisyDE** strategy (76.3% for **ChainDEwithMetaTuneRecentering** vs. 81.7% for **NoisyDE**) - and almost always outperforms the rescaling used in the original Nevergrad.

5 Conclusions and Future Work

We have investigated the scaling of the variance of random sampling in order to minimize the expected regret. While previous work [6] had already shown that, in the context of the sphere function, the optimal scaling factor is not identical to that of the prior distribution from which the optimum is sampled (unless the sample size is exponentially large in the dimension), it did not answer the question how to scale the variance optimally. We have proven that a standard deviation scaled as $\sigma = \sqrt{\log(\lambda)/d}$ gives, with probability at least $1/2$, a sample

that is significantly closer to the optimum than the previous known strategies. We have also proven that the gain achieved by our scaling strategy is asymptotically optimal and that any decent scaling factor is asymptotically at most as large as our suggestion.

The empirical assessment of our rescaled sampling strategy confirmed decent performance not only on the sphere function, but also on other classical benchmark problems. We have furthermore given indications that the sampling might help improve state-of-the-art numerical heuristics based on differential evolution or using Bayesian surrogate models. Our proposed one-shot method performs best in many cases, sometimes outperformed by e.g. `AvgLHS`, but is stable on a wide range of problems and meaningful also as an initialization method (as opposed to `AvgLHS`). Whereas our theoretical results can be extended to quadratic forms (by conservation of barycenters through linear transformations), an extension to wider families of functions (e.g., families of functions with order 2 Taylor expansion) is not straightforward. Apart from extending our results to broader function classes, another direction for future work comprises extensions to the multi-epoch case. Our empirical results on DE and BO gives a first indication that a properly scaled variance can also be beneficial in iterative sampling. Note, however, that in the latter case, we only adjusted the initialization, not the later sampling steps. This forms another promising direction for future work.

Acknowledgements. This work was initiated at Dagstuhl seminar 19431 on Theory of Randomized Optimization Heuristics.

References

1. Atanassov, E.I.: On the discrepancy of the Halton sequences. *Math. Balkanica (NS)* **18**(1–2), 15–32 (2004)
2. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **13**, 281–305 (2012)
3. Bossek, J., Doerr, C., Kerschke, P.: Initial design strategies and their effects on sequential model-based optimization. In: *Proceeding of the Genetic and Evolutionary Computation Conference (GECCO 2020)*. ACM (2020). <https://arxiv.org/abs/2003.13826>
4. Bossek, J., Kerschke, P., Neumann, A., Neumann, F., Doerr, C.: One-shot decision-making with and without surrogates. *CoRR abs/1912.08956* (2019). <http://arxiv.org/abs/1912.08956>
5. Bubeck, S., Munos, R., Stoltz, G.: Pure exploration in multi-armed bandits problems. In: *Gavaldà, R., Lugosi, G., Zeugmann, T., Zilles, S. (eds.) ALT 2009. LNCS (LNAI), vol. 5809, pp. 23–37*. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04414-4_7
6. Cauwet, M.L., et al.: Fully parallel hyperparameter search: reshaped space-filling. *arXiv preprint arXiv:1910.08406* (2019)
7. Dick, J., Pillichshammer, F.: *Digital Nets and Sequences*. Cambridge University Press, Cambridge (2010)
8. Ergezer, M., Sikder, I.: Survey of oppositional algorithms. In: *14th International Conference on Computer and Information Technology (ICCIT 2011)*, pp. 623–628 (2011)

9. Esmailzadeh, A., Rahnamayan, S.: Enhanced differential evolution using center-based sampling. In: 2011 IEEE Congress of Evolutionary Computation (CEC), pp. 2641–2648 (2011)
10. Esmailzadeh, A., Rahnamayan, S.: Center-point-based simulated annealing. In: 2012 25th IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), pp. 1–4 (2012)
11. Feurer, M., Springenberg, J.T., Hutter, F.: Initializing Bayesian hyperparameter optimization via meta-learning. In: AAAI (2015)
12. Halton, J.: On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.* **2**, 84–90 (1960). <http://eudml.org/doc/131448>
13. Hammersley, J.M.: Monte-Carlo methods for solving multivariate problems. *Ann. N. Y. Acad. Sci.* **86**(3), 844–874 (1960)
14. James, W., Stein, C.: Estimation with quadratic loss. In: Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Contributions to the Theory of Statistics, vol. 1, pp. 361–379. University of California Press (1961). <https://projecteuclid.org/euclid.bsmmsp/1200512173>
15. Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**(4), 455–492 (1998)
16. Maaranen, H., Miettinen, K., Mäkelä, M.: Quasi-random initial population for genetic algorithms. *Comput. Math. Appl.* **47**(12), 1885–1895 (2004)
17. Mahdavi, S., Rahnamayan, S., Deb, K.: Center-based initialization of cooperative co-evolutionary algorithm for large-scale optimization. In: 2016 IEEE Congress on Evolutionary Computation (CEC), pp. 3557–3565 (2016)
18. Matoušek, J.: *Geometric Discrepancy*, 2nd edn. Springer, Berlin (2010)
19. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**(2), 239–245 (1979)
20. Meunier, L., Doerr, C., Rapin, J., Teytaud, O.: Variance reduction for better sampling in continuous domains (2020)
21. Niederreiter, H.: *Random Number Generation and Quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia (1992)
22. Rahnamayan, S., Wang, G.G.: Center-based sampling for population-based algorithms. In: 2009 IEEE Congress on Evolutionary Computation, pp. 933–938, May 2009. <https://doi.org/10.1109/CEC.2009.4983045>
23. Rapin, J., Teytaud, O.: Nevergrad - a gradient-free optimization platform (2018). <https://GitHub.com/FacebookResearch/Nevergrad>
24. Stein, C.: Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In: Proceeding of the Third Berkeley Symposium on Mathematical Statistics and Probability, Contributions to the Theory of Statistics, vol. 1, pp. 197–206. University of California Press (1956). <https://projecteuclid.org/euclid.bsmmsp/1200501656>
25. Storn, R., Price, K.: Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *J. Glob. Optim.* **11**(4), 341–359 (1997)
26. Surry, P.D., Radcliffe, N.J.: Inoculation to initialise evolutionary search. In: Fogarty, T.C. (ed.) AISB EC 1996. LNCS, vol. 1143, pp. 269–285. Springer, Heidelberg (1996). <https://doi.org/10.1007/BFb0032789>

27. Teytaud, O., Gelly, S., Mary, J.: On the ultimate convergence rates for isotropic algorithms and the best choices among various forms of isotropy. In: Runarsson, T.P., Beyer, H.-G., Burke, E., Merelo-Guervós, J.J., Whitley, L.D., Yao, X. (eds.) PPSN 2006. LNCS, vol. 4193, pp. 32–41. Springer, Heidelberg (2006). https://doi.org/10.1007/11844297_4
28. Yang, X., Cao, J., Li, K., Li, P.: Improved opposition-based biogeography optimization. In: The Fourth International Workshop on Advanced Computational Intelligence, pp. 642–647 (2011)
29. Zhang, A., Zhou, Y.: On the non-asymptotic and sharp lower tail bounds of random variables (2018)