Novelty-Based Restarts for Evolution Strategies

Giuseppe Cuccu, Faustino Gomez and Tobias Glasmachers
Dalle Molle Institute for Artificial Intelligence (IDSIA)
Università della Svizzera Italiana, 6928 Manno-Lugano, Switzerland
Email: {giuse,tino,tobias}@idsia.ch

Abstract—**A major limitation in applying evolution strategies to black box optimization is the possibility of convergence into bad local optima. Many techniques address this problem, mostly through restarting the search. However, deciding the new start location is nontrivial since neither a good location nor a good scale for sampling a random restart position are known. A black box search algorithm can nonetheless obtain some information about this location and scale from past exploration. The method proposed here makes explicit use of such experience, through the construction of an archive of *novel solutions* during the run. Upon convergence, the most "novel" individual found so far is used to position the new start in the least explored region of the search space, actively looking for a new basin of attraction. We demonstrate the working principle of the method on two multi-modal test problems.**

keywords: evolution strategies, novelty search, restart strategies, black-box optimization

I. INTRODUCTION

Modern Evolution Strategies (ESs; [3, 5, 15]) are powerful "black-box" optimization algorithms that efficiently self-adapt their search strategy (i.e. the mutation step sizes) to local properties of the search space, and are invariant to linear transformations and monotone scaling of the fitness function. These features ensure high performance even on ill-conditioned problems with steep optima and nearly flat fitness plateaus.

However, the price to be payed for following the fitness gradient efficiently at all scales is that most ESs converge to a local minimum, after which no further progress can be made because the mutation step sizes vanish. Therefore, additional mechanisms are needed to continue the search in multi-modal fitness landscapes. Restart strategies [1, 2, 6] offer a simple and efficient mechanism for handling multi-modality that can be easily applied to any evolutionary search algorithm.

Recently, Novelty Search (NS; [12, 13]) has been proposed as a way to sustain evolutionary progress indefinitely. Unlike standard evolutionary search, where selection is based on a fixed objective or fitness function, in NS individuals are selected based on their "novelty" with respect to the solutions seen so far. The hope is that the process of continually searching for novel individuals will lead to increasingly complex solutions that, as a side-effect, eventually solve problems of interest.

While this goal-independent approach is conceptually interesting, its real potential for optimization seems to be in maintaining diversity in the context of conventional evolutionary search (i.e. fitness-based selection) [4].

In this paper, ideas from novelty search are used to provide a principled restart strategy that allows ES to escape local minima and continue making progress, while preserving scale invariance. The approach is tested on two multi-modal test problem: (1) a scaleable objective function where the length of a fitness plateaus and the number of local optima can be adjusted, and (2) the standard Rastrigin benchmark.

II. NATURAL EVOLUTION STRATEGIES

The most notable strengths of Evolution Strategies (ES) are fast self-adaptation of the search distribution to the local properties of the fitness landscape, and invariance of the search to transformations of the search space–importantly, translation and scaling. Modern ES, such as CMA-ES [10] and xNES [8], are even invariant under arbitrary affine linear transformations of the search space (of course, up to initialization). And, rank-based fitness assignment makes ES invariant under monotone transformations of the fitness function.

This study relies on the modern and flexible class of *Natural Evolution Strategies* (NES; [7, 8, 16–19]). Both hill-climbers and population-based variants of NES have been derived from the simple and powerful principle of adapting a parameterized search distribution in order to optimize the expected fitness by means of natural gradient descent. This general paradigm can be applied to all kinds of search distributions. Gaussians with different subsets of adaptive parameters have been treated in the literature, such as adaptation of the full covariance matrix [7, 8, 17–19], and diagonal covariance matrices [16].

Here we are not striving to improve the performance of a particular ES, and the approach is not tailored or limited to a particular algorithm. The aim is to avoid bias, which is best achieved by using as basic an ES as possible. To this end, in the following a natural evolution strategy is derived which uses a radial Gaussian, the simplest type of search distribution.

In each generation, the population-based NES algorithm samples $\lambda \in \mathbb{N}$ individuals $\mathbf{z}_k \sim \mathcal{N}(\mathbf{z} \,|\, \theta)$, $k \in \{1, \dots, \lambda\}$, i.i.d. from its Gaussian search distribution, parameterized by $\theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with the goal of minimizing a fitness function $f : \mathbb{R}^d \to \mathbb{R}$. Let $p(\mathbf{z} \,|\, \theta)$ denote the density of the Gaussian with parameters θ. Then, the expected fitness under the search distribution is

$$J(\theta) = \mathbb{E}_\theta[f(\mathbf{z})] = \int f(\mathbf{z})\; p(\mathbf{z} \,|\, \theta)\; d\mathbf{z} \ .$$

The gradient w.r.t. the parameters can be rewritten as

$$\begin{aligned}\nabla_\theta J(\theta) &= \nabla_\theta \int f(\mathbf{z})\ p(\mathbf{z}\,|\,\theta)\ d\mathbf{z} \\ &= \mathbb{E}_\theta\left[f(\mathbf{z})\ \nabla_\theta \log\left(p(\mathbf{z}\,|\,\theta)\right)\right]\ \ ,\end{aligned}$$

(see [19] for the full derivation) from which we obtain the Monte Carlo estimate

$$\nabla_\theta J(\theta) \approx \frac{1}{\lambda}\sum_{k=1}^{\lambda} f(\mathbf{z}_k)\ \nabla_\theta \log\left(p(\mathbf{z}_k\,|\,\theta)\right) \tag{1}$$

of the search gradient. The key step then consists of replacing this gradient, pointing into the direction of (locally) steepest descent w.r.t. the given parameterization, by the natural gradient

$$\widetilde{\nabla}_\theta J = \mathbf{F}^{-1}\nabla_\theta J(\theta)\ \ ,$$

where $\mathbf{F} = \mathbb{E}\left[\nabla_\theta \log\left(p\left(\mathbf{z}|\theta\right)\right) \nabla_\theta \log\left(p\left(\mathbf{z}|\theta\right)\right)^\top\right]$ is the Fisher information matrix; leading to a straightforward scheme of natural gradient descent for iteratively updating the search distribution

$$\theta \leftarrow \theta - \eta\widetilde{\nabla}_\theta J = \theta - \eta\mathbf{F}^{-1}\nabla_\theta J(\theta)\ \ ,$$

with learning rate parameter η. The sequence of (1) sampling an offspring population, (2) computing the corresponding Monte Carlo estimate of the fitness gradient, (3) transforming it into the natural gradient, and (4) updating the search distribution, constitutes one generation of NES.

In order to render the algorithm invariant under monotonic (rank preserving) transformations of the fitness values, *fitness shaping* [19] is used to normalize the fitness into rank-based *utilities* $u_k \in \mathbb{R}$, $k \in \{1, \ldots, \lambda\}$. The individuals are ordered by fitness, with $\mathbf{z}_{1:\lambda}$ and $\mathbf{z}_{\lambda:\lambda}$ denoting the most and least fit offspring, respectively. The distribution parameters are then updated using the "fitness-shaped" gradient:

$$\nabla_\theta J = \sum_{k=1}^{\lambda} u_k \cdot \nabla_{(\theta)} \log\left(p(\mathbf{z}_{k:\lambda}\,|\,\theta)\right)\ \ . \tag{2}$$

Typically, the utility values are either non-negative numbers that sum to one, or a shifted variant with zero mean.

An efficient scheme for Gaussian distributions with fully adaptive covariance matrix (called xNES) has been derived in [8]. Here, we restrict the search distribution to have the same variance in all dimensions, with parameters $\theta = (\boldsymbol{\mu}, \sigma)$, where $\boldsymbol{\mu}$ and σ are the mean and standard deviation of the Gaussian density

$$p(\mathbf{z}\,|\,\theta) = \frac{1}{(\sqrt{2\pi}\sigma)^d} \cdot \exp\left(-\frac{\|\mathbf{z}-\boldsymbol{\mu}\|^2}{2\sigma^2}\right)\ \ .$$

Using the same exponential local coordinates as in [8], we arrive at the update equations

$$\begin{aligned}\boldsymbol{\mu} &\leftarrow \boldsymbol{\mu} + \eta_\mu \cdot \sigma \cdot \sum_{i=1}^{\lambda} u_k \cdot \mathbf{s}_k \\ \sigma &\leftarrow \sigma \cdot \exp\left(\frac{\eta_\sigma}{2} \cdot \sum_{i=1}^{\lambda} u_k \cdot \left(\|\mathbf{s}_k\|^2 - d\right)\right)\ \ ,\end{aligned}$$

where $\mathbf{s}_k \sim \mathcal{N}(0, I)$ are standard normally distributed random variables used to generate the offspring $\mathbf{z}_k = \boldsymbol{\mu} + \sigma \cdot \mathbf{s}_k$. This version of NES with radial search distributions is referred to as *radial NES* (rNES).

III. NOVELTY SEARCH

In [13], Lehman and Stanley present the idea of *Novelty Search* where individuals in an evolving population are selected based solely on how different they are compared to all of the other solutions evaluated so far. Each individual, x, in the population is assigned a novelty score that is computed as the average *novelty distance* from its k-nearest neighbors in both the population and an archive:

$$\text{nov}(x_i) = \frac{1}{k}\sum_{j=1}^{k} dist(x_i, x_{i:j}), \tag{3}$$

where $x_{i:j}$ is the j-th nearest neighbor of x_i with respect to novelty distance $dist(\cdot,\cdot)$. The particular choice of distance measure is user-defined and problem specific. For example, novelty could be defined as simply as the Euclidean distance ($dist(x, y) = \|x - y\|$) between genotypes. In a more complex setting typically found in neuroevolution, controllers are generated from the genotype, and evaluated in a sequential decision task. Novelty is then computed on observation-action sequences.

Initially, the archive is empty. When a new individual is created, if its novelty exceeds the *novelty threshold*, then it is inserted into the archive. This threshold is subject to a self-adaptation mechanism with the goal to keep the number of individuals placed in the archive per generation around some fixed fraction (see [4] for details).

Notice there is no notion of "fitness" in the normal sense—the probability of being selected is not determined by a fixed objective function, but instead depends entirely on the current state (population) and history (archive) of a given evolutionary run.

The archive provides a memory of previously seen novel behaviors that may no longer be present in the population, but does not preclude the evaluation of non-novel solutions as in e.g. tabu search [9]. Because solutions that are novel (with respect to the current knowledge of the system) are selected for, their offspring have a better chance of being novel themselves. This, in turn, diverts the search away from wastefully retracing the same ground.

An interesting aspect of the novelty signal is that it is non-stationary. Initially points in an unexplored region of the search space receive high novelty scores, but as the evolutionary algorithm explores the space and places more and more individuals from this region into the archive, they become less and less novel. Thus, the search algorithm will tend to focus on other less explored parts of the search space. In a bounded search space the expected behavior is that the archive will gradually fill with representative solutions, reducing the average novelty of the whole space over time, resulting in adaptive sampling of the space at higher and higher resolutions.

IV. NOVELTY IN BLACK BOX-TYPE OPTIMIZATION

In this section we investigate how a novelty-based archive can be used to improve the performance of evolutionary strategies on multi-modal fitness landscapes while preserving the properties that make these algorithms attractive for black box search. Ideally, the resulting scheme should exhibit the same invariance properties as the underlying evolution strategy, and should introduce as few as possible additional parameters that need problem-specific tuning. The method is described with the rNES algorithm in mind, however, most steps are generic and can easily be adapted to other evolution strategies.

Novelty search has mostly been applied in neuroevolution tasks. The biggest difference to novelty search in black-box search is the choice of the novelty score. By definition, in a black box setting a behavioral or otherwise task-specific distance measure is not available. Thus, novelty necessarily needs to be measured based on the genotype.

The next two subsections (IV-A,IV-B) analyze two possible ways of combining novelty with fitness, and provide the rationale for using novelty as a basis for restarting an ES, which is present in the last subsection (IV-C).

A. Blending Fitness and Novelty

In [4] fitness and novelty are combined as follows:

$$\rho \cdot n(\mathbf{z}_k) + (1-\rho) \cdot f(\mathbf{z}_k)$$

with $n(\cdot)$ denoting novelty (w.r.t. the archive) and $f(\cdot)$ fitness. For NES this expression would replace $f(\mathbf{z}_k)$ in equation (1). This formula allows for smooth blending between exploration of the fitness landscape and exploitation of fitness by following the global trend and descending into a local optimum.

In order to mix fitness and novelty meaningfully, the two measures must be normalized. It is possible to normalize over the archive or over a sliding time window. However, rank-based selection, as found e.g. in the NES algorithms presented in section II, points to a principled solution: instead of combining raw fitness and novelty values, first rank the population independently according to both measures, and then blend the corresponding (already normalized) utility values by replacing u_k in equation (2) with:

$$\rho \cdot u_k^n + (1-\rho) \cdot u_k^f \quad ,$$

where u_k^n is the utility of individual $\mathbf{z}_k$ with respect to novelty (equation 3), u_k^f with respect to fitness, and $\rho \in [0,1]$ is the blending parameter. Alternatively, ranks can be blended, which requires minimal changes in the way the utility values are computed.

It turns out that despite its conceptual simplicity, blending does not result in the desired effect because of the scale invariance of the algorithm. The problem is best explained with a scale-invariant fitness function like the sphere function, $f(x) = \|x\|^2$: depending on the value of ρ, there is either a trend to shrink or to widen the search distribution, and this trend is the same on all scales. Thus, there is a sudden shift from convergence into the optimum (small ρ) to exploration by means of divergence (large ρ), rather than a smooth transition.

B. Switching between Fitness and Novelty

The above observation suggests restricting ρ to $\{0,1\}$, corresponding to a hard switch between fitness and novelty. This strategy is arguably even closer to the idea of exploring the most promising regions in search space, amounting to two alternating phases: in the fitness-based phase of the search the evolution strategy identifies a (local) optimum. Once convergence or stagnation are detected, the algorithm switches to novelty mode in order to escape the local optimum or plateau, and hopefully reach a different, ultimately better attractor basin. In the next fitness-based phase the ES converges to the next local optimum, and the process starts over.

In fitness mode the evolution strategy operates normally, until its stopping criterion is met. In black-box search, this usually amounts to detecting convergence; however, this stopping criterion necessarily violates the requirement of scale invariance, as the algorithm would otherwise just keep making insignificant improvements forever. The same holds for the switching mechanism from fitness to novelty mode. Thus, it is canonical to use a standard stopping criterion here.

The rNES algorithm provides us with a particularly easy criterion to check, as σ, is an explicit step size. Once σ falls below some small threshold (e.g. as 10^{-5}, or the limit of numerical accuracy), convergence is considered to be achieved, the fitness phase stops, and the novelty-based restart method looks for the next restart position.

C. Novelty Restarts

Notice that, since novelty here is based on genotypic similarity, the choice of $\rho = 1$ means that the evaluation of the fitness function can be skipped completely. While this may mean a big savings in terms of sample complexity, how can novelty alone be used to decide when to switched back to fitness? How far away from the last optima is far enough in to escape the local minimum? The distribution variance, σ is no longer a useful criterion as novelty search will cause the search to diverge. And, even if the ES escapes the previous local minimum, novelty will then only delay convergence to the next minimum. Furthermore, it seems dangerous to add individuals to the novelty archive without evaluating their fitness because it allows regions of the search space to be marked as 'explored'—the global optimum may be overlooked and, even worse, actively ignored.

This insight suggests that instead of actually running the novelty search, one should just use the information gathered so far in the novelty archive to restart the search with a new $\boldsymbol{\mu}$ and σ. The goal is to continue the search in a less explored region of the space. The approach proposed here is to select the individual from the archive with the highest novelty score as the new $\boldsymbol{\mu}$ of the search distribution, and use the distance from this individual to the rest of the archive (e.g. the mean distance to its 5 nearest neighbors), as the new standard deviation σ. Note how this scheme, except for the stopping criterion, preserves all invariance properties of the underlying evolution strategy. The new mean, being the most novel point in archive, is by definition the most isolated of

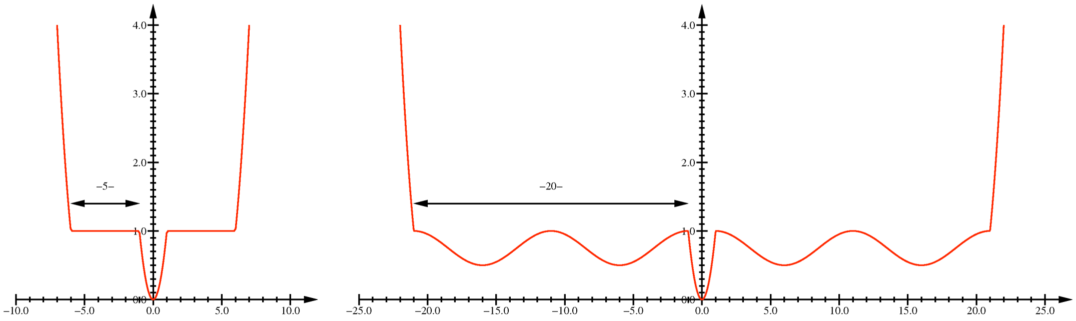

Fig. 1. The function $f_{\ell,w}$ for the configurations ($\ell = 5$, $w = 0$) and ($\ell = 20$, $w = 2$). In the first case the difficulty for the evolution strategy is to cross the (possibly long) fitness plateau; in the much harder second case there are deceptive local minima.

	$\ell = 5$	$\ell = 15$	$\ell = 30$
$w = 0$	97	87	85
$w = 2$	19	3	0
$w = 10$	60	13	2

TABLE I

PERCENTAGE OF TRIALS IN WHICH RNES *without novelty restarts* FINDS THE GLOBAL OPTIMUM OF $f_{\ell,w}$, OUT OF 100 RUNS.

	$\ell = 5$	$\ell = 15$	$\ell = 30$
$w = 0$	79	298	319
$w = 2$	267	668	(677)*
$w = 10$	101	792	903

TABLE II

MEDIAN NUMBER OF GENERATIONS (OVER 5 RUNS) REQUIRED FOR RNES *with novelty restarts* TO IDENTIFY THE GLOBAL OPTIMUM. THE MOST DIFFICULT CONFIGURATION (PLATEAU LENGTH $\ell = 30$ WITH ONLY $w = 2$ LOCAL OPTIMA PER SIDE) WAS SOLVED IN ONE OUT OF FIVE RUNS, AFTER 667 GENERATIONS.

the search points visited so far, and therefore far from where the ES converged before the restart. Because this point was generate by the ES it on the same scale. Likewise, the new σ is proportional to the scale of the points seen so far. Once the ES restarts from the new point, the novelty at that point will decrease such that subsequent restarts will avoid re-exploring the same region.

By only using novelty to build the archive and not to update the search distribution (i.e. novelty search), this approach is related to other evolutionary methods that use an archive, typically for multi-objective optimization [11, 20], except that here the archive is based on the concept of novelty.

V. EXPERIMENTAL EVALUATION

The following experiments are designs to test the working principle of novelty-based restarts, and to study how it scales with problem difficulty. By construction restarts neither help nor harm on unimodal fitness landscapes where it is up to the underlying ES to identify the optimum. Therefore the experiments focus on fitness function that exhibit plateaus and multi-modality.

As in [16], the population size was set to $\lambda = 4 + \lfloor 3\log(d) \rfloor$, and the learning rates to $\eta_{\mu} = 1$, and $\eta_{\sigma} = \frac{(3+\log(d))}{5\sqrt{d}}$, for all experiments.

A. Plateaus and Local Minima

The first function tests the ability of the algorithm to cross long flat fitness plateaus and to identify the global optimum in a deceptive landscape of much larger attractor basins with local minima. For this purpose we introduce the scaleable fitness function

$$f_{\ell,w}(x) = \begin{cases} x^2 & \text{for } |x| \leq 1 \\ 1 - \frac{1}{2}\sin^2\left(\frac{\pi w(|x|-1)}{\ell}\right) & \text{for } 1 \leq |x| \leq \ell + 1 \\ (|x| - \ell)^2 & \text{for } |x| \geq \ell + 1 \end{cases}$$

with parameters $\ell \geq 0$ and $w \in \mathbb{N}$. This function is symmetric around its global optimum at 0, and based on a standard parabola. However, the parabola is chopped open at $+1$ and -1, with its branches shifted outwards by the plateau length ℓ. Additionally, w waves or local optima can be added to each side of the plateau (figure 1). The search is started at $\boldsymbol{\mu}_0 = \ell + 1$ with a standard deviation of $\sigma = 1$.

rNES was run with and without novelty restarts on this function, with all combinations of $\ell \in \{5, 15, 30\}$ and $w \in \{0, 2, 10\}$, and a budget of 10,000 generations (corresponding to 40,000 fitness evaluations). The stopping criterion, which was also used to trigger restarts, was $\sigma < 10^{-10}$.

Table I shows the percentage of runs in which rNES (without restarts) converged into the global optimum. By itself. rNES copes well only with plateaus of limited length, while local optima are indeed deceptive. The task is hardest for only two local optima per side because of the size of the attractor basins.

Table II shows the median number of generations it took rNES with novelty restarts to identify the optimum with an accuracy of 10^{-5}. In contrast, the novelty restart strategy manages to identify the optimum reliably, and scales gracefully with respect to both plateau length and number of local optima. It is clear from figure 2 that the archive adapts itself to the current scale of the search distribution and fills the search

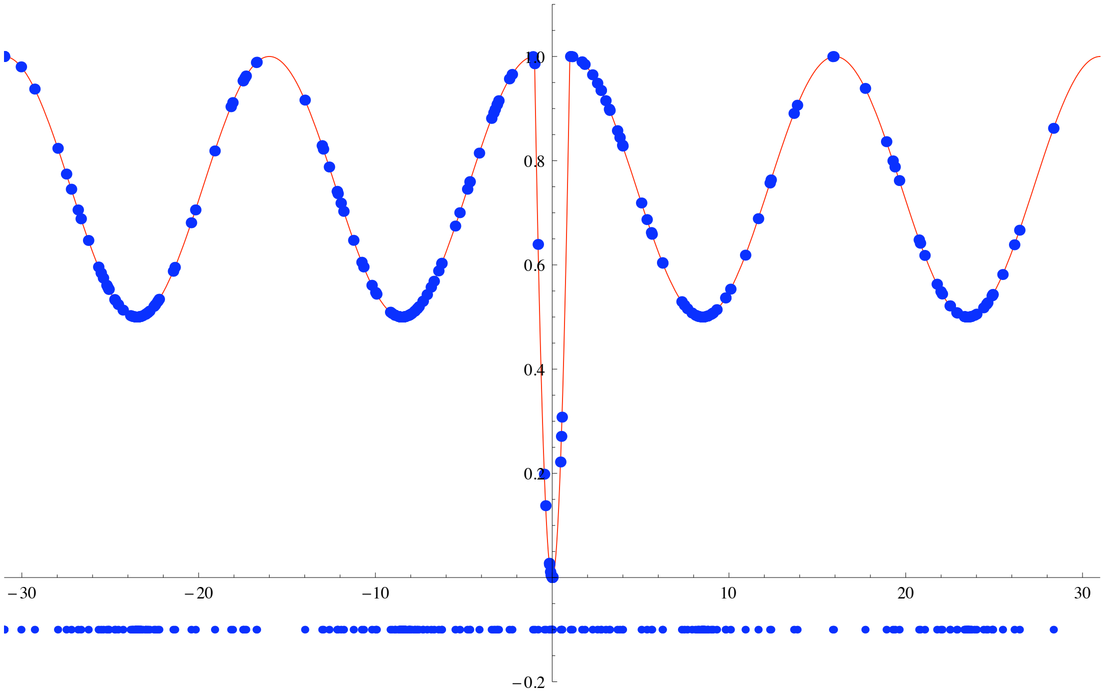

Fig. 2. Archive at the end of a typical run of the novelty restart algorithm for $\ell = 30$ and $w = 2$.

space until it finally finds the global optimum. However, the fitness-based component of the search makes sure that the interesting regions around the local optima are sampled at a much finer scale than the rest of the search space.

B. Scaling to Higher Dimensions

The function $f_{\ell,w}$ is a worst-case scenario for an evolution strategy. The lack of global trend makes it a continuous version of the needle-in-a-haystack problem, that requires a thorough search of the space in order to finally find the small attractor of the global optimum. The archive here helps a lot, as any element in it acts as a milestone, discouraging the search from wasting time in the close surrounding.

In contrast, the Rastrigin function

$$f_{\text{Rastrigin}} : \mathbb{R}^d \to \mathbb{R}, \quad x \mapsto 10d + \sum_{i=1}^{d} \left[x_i^2 - 10 \cdot \cos(2\pi x_i)\right]$$

has a quadratic global trend, overlayed with a grid of local optima. This benchmark is used to investigate the scaling of novelty-based restarts to higher-dimensional problems, where it is not feasible for the archive to cover the search space.

Each run was started by initializing $\boldsymbol{\mu}_0$ randomly from a uniform distribution over the interval $[-5, 5]^d$, with $\sigma = 1$. As in the previous experiment above, for the runs without restarts, the percentage of successful runs was recorded (Table III). For the restart runs the number of generations required to find the optimum is reported since almost all the runs were successful (Table III). Experiments were conducted for dimensions $d \in \{2, 5, 10, 20\}$, with a maximum of 10,000 generations for each run.

$d = 2$	$d = 5$	$d = 10$	$d = 20$
24.6	25.6	18.9	16.2

TABLE III
PERCENTAGE OF TRIALS IN WHICH RNES *without novelty restarts* FINDS THE GLOBAL OPTIMUM OF THE RASTRIGIN FUNCTION, OUT OF 1,000 RUNS.

$d = 2$	$d = 5$	$d = 10$	$d = 20$
352	234	64.3	1374*

TABLE IV
MEDIAN NUMBER OF GENERATIONS (OVER 25 RUNS) IT TAKES RNES *with novelty restarts* TO IDENTIFY THE GLOBAL OPTIMUM OF THE RASTRIGIN FUNCTION. IN DIMENSION $d = 20$ THE ALGORITHM FAILED TO IDENTIFY THE GLOBAL OPTIMUM IN ONLY 2 OUT OF 25 RUNS WITHIN ITS BUDGET OF 10,000 GENERATIONS.

The results show that novelty restarts scale well to relatively high dimensions. Even in the hypercube $[-1.5, 1.5]^d$ the Rastrigin function has 3^d local optima. If the algorithm would need to represent all of these (and even other) local optima in its archive, then a scaling of at least $\Omega(3^d)$ would be expected. In contrast, the results show a far more graceful scaling to higher dimensions. This is possible because the Rastrigin problem is much less deceptive than the purpose-made $f_{\ell,w}$ function.

VI. DISCUSSION AND CONCLUSION

We have investigated novelty-based search for evolutionary black-box optimization. The focus was on preserving the invariance properties that make evolution strategies powerful for these problems. When combining fitness and novelty under these constraints, our analysis revealed that novelty-based search phases can be understood as a restart strategy for fitness-based search.

In contrast to simpler restart strategies, these restarts are informed decisions taken based on information explicitly stored in an archive of already explored solutions. Our restart strategy naturally respects important invariance properties of the ES, such as translation, rotation, and scale invariance, since it is based solely on information produced by the evolution strategy itself.

A deceptive benchmark was designed to highlight the benefits of novelty-based restarts. However, the strategy of filling an archive with all local optima of the search space bears the risk of exponential sample complexity with search space dimension. Experiments on the less deceptive, and arguably more realistic benchmark defined by the Rastrigin function, showed that in practice such a archive-filling behavior is not necessary to solve the task, resulting in a good scalability with respect to search space dimensionality.

It remains unclear when exactly novelty-based restarts are better than other restart schemes. We believe that this question is best answered by a thorough experimental comparison, left for future work.

ACKNOWLEDGMENTS

This work was funded through the 7th framework program of the EU under grant numbers 231576 (STIFF project), and 228844 (NanoBioTouch project), and by Swiss National Science Foundation grant no. 120061 (EVOCMP).

REFERENCES

[1] A. Auger and N. Hansen. A restart CMA evolution strategy with increasing population size. In B. McKay et al., editors, *The 2005 IEEE International Congress on Evolutionary Computation (CEC'05)*, volume 2, pages 1769–1776, 2005.

[2] David Beasley, David R. Bull, and Ralph R. Martin. A Sequential Niche Technique for Multimodal Function Optimization. *Evolutionary Computation*, 1(2):101–125, 1993.

[3] H.-G. Beyer and H.-P. Schwefel. "Evolution strategies"–a comprehensive introduction. *Natural Computing*, 1:3–52, 2002.

[4] Giuseppe Cuccu and Faustino Gomez. When Novelty is not Enough. In *Proceedings of evo* 2011 (Torino, Italy)*, 2011.

[5] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, Berlin, Germany, 2003.

[6] Alex S. Fukunaga. Restart scheduling for genetic algorithms. In *Genetic Algorithms: Proceedings of the Seventh International Conference*, pages 357–358, 1998.

[7] T. Glasmachers, T. Schaul, and J. Schmidhuber. A Natural Evolution Strategy for Multi-Objective Optimization. In *Parallel Problem Solving from Nature (PPSN)*, 2010.

[8] T. Glasmachers, T. Schaul, Y. Sun, D. Wierstra, and J. Schmidhuber. Exponential Natural Evolution Strategies. In *Genetic and Evolutionary Computation Conference (GECCO)*, Portland, OR, 2010.

[9] F. Glover and R. Marti. Tabu search. *Metaheuristic Procedures for Training Neutral Networks*, pages 53–69, 2006.

[10] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.

[11] J. D. Knowles and D. W. Corne. The pareto archived evolution strategy : A new baseline algorithm for pareto multiobjective optimisation. In *Proceedings of the Congress on Evolutionary Computation (CEC'99)*, pages 98–105, 1999.

[12] Joel Lehman and Kenneth Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation Journal*, 2010.

[13] Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*. MIT, 2008.

[14] Joel Lehman and Kenneth O. Stanley. Efficiently evolving programs through the search for novelty. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-10)*, 2010.

[15] I. Rechenberg and M. Eigen. Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. 1973.

[16] T. Schaul, T. Glasmachers, and J. Schmidhuber. High Dimensions and Heavy Tails for Natural Evolution Strategies. under review, 2011.

[17] Y. Sun, D. Wierstra, T. Schaul, and J. Schmidhuber. Efficient Natural Evolution Strategies. In *Genetic and Evolutionary Computation Conference (GECCO)*, 2009.

[18] Y. Sun, D. Wierstra, T. Schaul, and J. Schmidhuber. Stochastic Search using the Natural Gradient. In *International Conference on Machine Learning (ICML)*, 2009.

[19] D. Wierstra, T. Schaul, J. Peters, and J. Schmidhuber. Natural Evolution Strategies. In *Proceedings of the Congress on Evolutionary Computation (CEC08), Hongkong*. IEEE Press, 2008.

[20] Jingqiao Zhang and Arthur C. Sanderson. Self-adaptive differential evolution with fast and reliable convergence performance. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC'07)*, pages 2251–2258, 2007.