



Covariance Matrix Adaptation MAP-Annealing

Matthew C. Fontaine
University of Southern California
Los Angeles, CA, USA
mfontain@usc.edu

Stefanos Nikolaidis
University of Southern California
Los Angeles, CA, USA
nikolaid@usc.edu

ABSTRACT

Single-objective optimization algorithms search for the single highest-quality solution with respect to an objective. Quality diversity (QD) optimization algorithms, such as Covariance Matrix Adaptation MAP-Elites (CMA-ME), search for a collection of solutions that are both high-quality with respect to an objective and diverse with respect to specified measure functions. However, CMA-ME suffers from three major limitations highlighted by the QD community: prematurely abandoning the objective in favor of exploration, struggling to explore flat objectives, and having poor performance for low-resolution archives. We propose a new quality diversity algorithm, Covariance Matrix Adaptation MAP-Annealing (CMA-MAE), that addresses all three limitations. We provide theoretical justifications for the new algorithm with respect to each limitation. Our theory informs our experiments, which support the theory and show that CMA-MAE achieves state-of-the-art performance and robustness.

ACM Reference Format:

Matthew C. Fontaine and Stefanos Nikolaidis. 2023. Covariance Matrix Adaptation MAP-Annealing. In *Genetic and Evolutionary Computation Conference (GECCO '23)*, July 15–19, 2023, Lisbon, Portugal. ACM, New York, NY, USA, 36 pages. <https://doi.org/10.1145/3583131.3590389>

1 INTRODUCTION

Consider an example problem of searching for celebrity faces in the latent space of a generative model. As a single-objective optimization problem, we specify an objective f that targets a celebrity such as Tom Cruise. A single-objective optimizer, such as CMA-ES [35], will converge to a single solution of high objective value, an image that looks like Tom Cruise as much as possible.

However, this objective has ambiguity. How old was Tom Cruise in the photo? Did we want the person in the image to have short or long hair? By instead framing the problem as a quality diversity optimization problem, we additionally specify a measure function m_1 that quantifies age and a measure function m_2 that quantifies hair length. A quality diversity algorithm [8, 66], such as CMA-ME [26], can then optimize for a collection of images that are diverse with respect to age and hair length, but all look like Tom Cruise.

While prior work [17, 20, 24, 26] has shown that CMA-ME solves such QD problems efficiently, three important limitations of the algorithm have been discovered. First, on difficult to optimize objectives, variants of CMA-ME will abandon the objective too soon [76], and instead favor exploring the measure space, the vector space defined

by the measure function outputs. Second, the CMA-ME algorithm struggles to explore flat objective functions [62]. Third, CMA-ME works well on high-resolution archives, but struggles to explore low-resolution archives [14, 21].¹

We propose a new algorithm, CMA-MAE, that addresses these three limitations. To address the first limitation, we derive an algorithm that smoothly blends between CMA-ES and CMA-ME. First, consider how CMA-ES and CMA-ME differ. At each step CMA-ES's objective ranking maximizes the objective function f by approximating the natural gradient of f at the current solution point [1]. In contrast, CMA-ME's improvement ranking moves in the direction of the natural gradient of $f - f_A$ at the current solution point, where f_A is a discount function equal to the objective of the best solution so far that has the same measure values as the current solution point. The function $f - f_A$ quantifies the gap between a candidate solution and the best solution so far at the candidate solution's position in measure space.

Our key insight is to *anneal the function f_A by a learning rate α* . We observe that when $\alpha = 0$, then our discount function f_A never increases and our algorithm behaves like CMA-ES. However, when $\alpha = 1$, then our discount function always maintains the best solution for each region in measure space and our algorithm behaves like CMA-ME. For $0 < \alpha < 1$, CMA-MAE smoothly blends between the two algorithms' behavior, allowing for an algorithm that spends more time on the optimization of f before transitioning to exploration. Figure 1 is an illustrative example of varying the learning rate α .

Our proposed annealing method naturally addresses the flat objective limitation. Observe that both CMA-ES and CMA-ME struggle on flat objectives f as the natural gradient becomes 0 in this case and each algorithm will restart. However, we show that, when CMA-MAE optimizes $f - f_A$ for $0 < \alpha < 1$, the algorithm becomes a descent method on the density histogram defined by the archive.

Finally, CMA-ME's poor performance on low resolution archives is likely caused by the non-stationary objective $f - f_A$ changing too quickly for the adaptation mechanism to keep up. Our archive learning rate α controls how *quickly* $f - f_A$ changes. We derive a conversion formula for α that allows us to derive equivalent α for different archive resolutions. Our conversion formula guarantees that CMA-MAE is the first QD algorithm invariant to archive resolution.

While our new annealing method benefits CMA-ME, our approach is also compatible with CMA-ME's differentiable quality diversity (DQD) counterpart, CMA-MEGA [21]. We apply the same modification in the DQD setting to form CMA-MAEGA. To evaluate CMA-MAEGA, we improve the latent space illumination (LSI) [21] domain by introducing a higher-dimensional domain based on StyleGAN2, capable of producing higher quality images. This new domain highlights the advantages of DQD in high-dimensional spaces and demonstrates the performance benefits of our annealing method.

¹We note that archive resolution affects the performance of all current QD algorithms.



This work is licensed under a Creative Commons Attribution International 4.0 License.
GECCO '23, July 15–19, 2023, Lisbon, Portugal
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0119-1/23/07.
<https://doi.org/10.1145/3583131.3590389>




Figure 1: An example of how different α values affect the function $f - f_A$ optimized by CMA-MAE after a fixed number of iterations. Here f is a bimodal objective where mode X is harder to optimize than mode Y , requiring more optimization steps, and modes X and Y are separated by measure m_1 . For $\alpha = 0$, the objective f is equivalent to $f - f_A$, as f_A remains constant. For larger values of α , CMA-MAE discounts region Y in favor of prioritizing the optimization of region X .

Overall, our work shows how a simple algorithmic change to CMA-ME addresses all three major limitations affecting CMA-ME’s performance and robustness. Our theoretical findings justify the aforementioned properties and inform our experiments, which show that CMA-MAE outperforms state-of-the-art QD algorithms and maintains robust performance across different archive resolutions.

2 PROBLEM DEFINITION

Quality Diversity. We adopt the quality diversity (QD) problem definition from prior work [21]. A QD problem consists of an objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that maps n -dimensional solution parameters to the scalar value representing the quality of the solution and k measure functions $m_i : \mathbb{R}^n \rightarrow \mathbb{R}$ or, as a vector function, $\mathbf{m} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ that quantify the behavior or attributes of each solution.² The range of $\mathbf{m}$ forms a measure space $S = \mathbf{m}(\mathbb{R}^n)$. The QD objective is to find a set of solutions $\theta \in \mathbb{R}^n$, such that $\mathbf{m}(\theta) = s$ for each s in S and $f(\theta)$ is maximized.

The measure space S is continuous, but solving algorithms need to produce a finite collection of solutions. Therefore, QD algorithms in the MAP-Elites [15, 59] family relax the QD objective by discretizing the space S . Given T as the tessellation of S into M cells, the QD objective becomes to find a solution θ_i for each of the $i \in \{1, \dots, M\}$ cells, such that each θ_i maps to the cell corresponding to $\mathbf{m}(\theta_i)$ in the tessellation T . The QD objective then becomes maximizing the objective value $f(\theta_i)$ of all cells:

$$\max \sum_{i=1}^M f(\theta_i) \quad (1)$$

The *differentiable quality diversity* (DQD) problem [21] is a special case of the QD problem where both the objective f and measures m_i are *first-order differentiable*.

3 PRELIMINARIES

We present several QD algorithms that solve derivative-free QD problems to provide context for our proposed CMA-MAE algorithm. Appendix D contains information about the DQD algorithm CMA-MEGA, which solves problems where the gradients of the objective and measure functions are available.

²In agent-based settings, such as reinforcement learning, the measure functions are sometimes called behavior functions and the outputs of each measure function are called behavioral characteristics or behavior descriptors.

MAP-Elites and MAP-Elites (line). The MAP-Elites QD algorithm produces an archive of solutions, where each cell in the archive corresponds to the provided tessellation T in the QD problem definition. The algorithm initializes the archive by sampling solutions from the solution space $\mathbb{R}^n$ from a fixed distribution. After initialization, MAP-Elites produces new solutions by selecting occupied cells uniformly at random and perturbing them with isotropic Gaussian noise: $\theta' = \theta_i + \sigma \mathcal{N}(0, I)$. For each new candidate solution θ' , the algorithm computes an objective $f(\theta')$ and measures $\mathbf{m}(\theta')$. MAP-Elites places θ' into the archive if the cell corresponding to $\mathbf{m}(\theta')$ is empty or θ' obtains a better objective value $f(\theta')$ than the current occupant. The MAP-Elites algorithm results in an archive of solutions that are diverse with respect to the measure function $\mathbf{m}$, but also high quality with respect to the objective f . Prior work [79] proposed the MAP-Elites (line) algorithm by augmenting the isotropic Gaussian perturbation with a linear interpolation between two solutions θ_i and θ_j : $\theta' = \theta_i + \sigma_1 \mathcal{N}(0, I) + \sigma_2 \mathcal{N}(0, 1)(\theta_i - \theta_j)$. **CMA-ME.** Covariance Matrix Adaptation MAP-Elites [26] combines the archiving mechanisms of MAP-Elites with the adaptation mechanisms of CMA-ES [35]. Instead of perturbing archive solutions with Gaussian noise, CMA-ME maintains a multivariate Gaussian of search directions $\mathcal{N}(0, \Sigma)$ and a search point $\theta \in \mathbb{R}^n$. The algorithm updates the archive by sampling λ solutions around the current search point $\theta_i \sim \mathcal{N}(\theta, \Sigma)$. After updating the archive, CMA-ME ranks solutions via a two stage ranking. Solutions that discover a new cell are ranked by the objective $\Delta_i = f(\theta_i)$, and solutions that map to an occupied cell e are ranked by the improvement over the incumbent solution θ_e in that cell: $\Delta_i = f(\theta_i) - f(\theta_e)$. CMA-ME prioritizes exploration by ranking all solutions that discover a new cell before all solutions that improve upon an existing cell. Finally, CMA-ME moves θ towards the largest improvement in the archive, according to the CMA-ES update rules. Prior work [21] showed that the improvement ranking of CMA-ME approximates a natural gradient of a modified QD objective (see Eq. 1).

4 PROPOSED ALGORITHMS

We present the CMA-MAE algorithm. While we focus on CMA-MAE, the same augmentations apply to CMA-MEGA to form the novel CMA-MAEGA algorithm (see Appendix D).

CMA-MAE. CMA-MAE is an algorithm that adjusts the rate the non-stationary QD objective $f - f_A$ changes. First, consider at a




Figure 2: Our proposed CMA-MAE algorithm smoothly blends between the behavior of CMA-ES and CMA-ME via an archive learning rate α . Each heatmap visualizes an archive of solutions across a 2D measure space, where the color of each cell represents the objective value of the solution.

high level how CMA-ME explores the measure space and discovers high quality solutions. The CMA-ME algorithm maintains a solution point θ and an archive A with previously discovered solutions. When CMA-ME samples a new solution θ' , the algorithm computes the solution's objective value $f(\theta')$ and maps the solution to a cell e in the archive based on the measure $\mathbf{m}(\theta')$. CMA-ME then computes the improvement of the objective value $f(\theta')$ of the new solution, over a discount function $f_A : \mathbb{R}^n \rightarrow \mathbb{R}$. In CMA-ME, we define $f_A(\theta')$ by computing the cell e in the archive corresponding to $\mathbf{m}(\theta')$ and letting $f_A(\theta') = f(\theta_e)$, where θ_e is the incumbent solution of cell e . The algorithm ranks candidate solutions by improvement $f(\theta') - f_A(\theta') = f(\theta') - f(\theta_e)$ and moves the search in the direction of higher ranked solutions.

Assume that CMA-ME samples a new solution θ' with a high objective value of $f(\theta') = 99$. If the current occupant θ_e of the corresponding cell has a low objective value of $f(\theta_e) = 0.3$, then the improvement in the archive $\Delta = f(\theta') - f(\theta_e) = 98.7$ is high and the algorithm will move the search point θ towards θ' . Now, assume that in the next iteration the algorithm discovers a new solution θ'' with objective value $f(\theta'') = 100$ that maps to the same cell as θ' . The improvement then is $\Delta = f(\theta'') - f(\theta') = 1$ as θ' replaced θ_e in the archive in the previous iteration. CMA-ME would likely move θ away from θ'' as the solution resulted in low improvement. In contrast, CMA-ES would move towards θ'' as it ranks only by the objective f , ignoring previously discovered solutions with similar measure values.

In the above example, CMA-ME moves away from a high performing region in order to maximize how the archive changes. However, in domains with hard-to-optimize objective functions, it is beneficial to perform more optimization steps towards the objective f before leaving each high-performing region [76].

Like CMA-ME, CMA-MAE maintains a discount function $f_A(\theta')$ and ranks solutions by improvement $f(\theta') - f_A(\theta')$. However, instead of maintaining an *elitist archive* by setting $f_A(\theta')$ equal to $f(\theta_e)$, we maintain a *soft archive* by setting $f_A(\theta')$ equal to t_e , where t_e is an acceptance threshold maintained for each cell in the archive A . When adding a candidate solution to the archive, we control the rate that t_e changes by the archive learning rate α as follows: $t_e \leftarrow (1 - \alpha)t_e + \alpha f(\theta')$.

The archive learning rate α in CMA-MAE allows us to control how quickly we leave a high-performing region of measure space. For example, consider discovering solutions in the same cell with objective value 100 in 5 consecutive iterations. The improvement values computed by CMA-ME against the elitist archive would be 100, 0, 0, 0, 0, thus CMA-ME would move rapidly away from this cell. The improvement values computed against the soft archive of CMA-MAE with $\alpha = 0.5$ would diminish smoothly as follows: 100, 50, 25, 12.5, 6.25, enabling further exploitation of the high-performing region.

Next, we walk through the CMA-MAE algorithm step-by-step. Algorithm 1 shows the pseudo-code for CMA-MAE with the differences from CMA-ME highlighted in yellow. First, on line 3 we initialize the acceptance threshold to $\min_f$. In each iteration we sample λ solutions around the current search point θ (line 6). For each candidate solution θ_i , we evaluate the solution and compute the objective value $f(\theta_i)$ and measure values $\mathbf{m}(\theta_i)$ (line 7). Next, we compute the cell e in the archive that corresponds to the measure values and the improvement Δ_i over the current threshold t_e (lines 8-9). If the objective crosses the acceptance threshold t_e , we replace the incumbent θ_e in the archive and increase the acceptance threshold t_e (lines 10-12). Next, we rank all candidate solutions θ_i by their improvement Δ_i . Finally, we step our search point θ and adapt our covariance matrix Σ towards the direction of largest improvement (lines 15-16) according to CMA-ES's update rules [35].

CMA-MAEGA. We note that our augmentations to the CMA-ME algorithm only affects how we replace solutions in the archive and how we calculate Δ_i . CMA-ME and CMA-MEGA replace solutions and calculate Δ_i identically, so we apply the same augmentations to CMA-MEGA to form a new DQD algorithm, CMA-MAEGA, in Appendix D.

5 THEORETICAL PROPERTIES OF CMA-MAE

We provide insights about the behavior of CMA-MAE for different α values. We include all proofs in Appendix E. CMA-MAEGA has similar theoretical properties discussed in Appendix F.

THEOREM 5.1. *The CMA-ES algorithm is equivalent to CMA-MAE when $\alpha = 0$, if CMA-ES restarts from an archive solution.*

```

1 CMA-MAE (evaluate,  $\theta_0$ ,  $N$ ,  $\lambda$ ,  $\sigma$ ,  $\min_f$ ,  $\alpha$ )
   input : An evaluation function evaluate that computes
           the objective and measures, an initial solution  $\theta_0$ ,
           a desired number of iterations  $N$ , a branching
           population size  $\lambda$ , an initial step size  $\sigma$ , a minimal
           acceptable solution quality  $\min_f$ , and an archive
           learning rate  $\alpha$ .
   result : Generate  $N\lambda$  solutions storing elites in an archive
           A.
2   Initialize solution parameters  $\theta$  to  $\theta_0$ , CMA-ES
   parameters  $\Sigma = \sigma I$  and  $p$ , where we let  $p$  be the
   CMA-ES internal parameters.
3   Initialize the archive A and the acceptance threshold  $t_e$ 
   with  $\min_f$  for each cell  $e$ .
4   for  $iter \leftarrow 1$  to  $N$  do
5     for  $i \leftarrow 1$  to  $\lambda$  do
6        $\theta_i \sim \mathcal{N}(\theta, \Sigma)$ 
7        $f, m \leftarrow \text{evaluate}(\theta_i)$ 
8        $e \leftarrow \text{calculate\_cell}(A, m)$ 
9        $\Delta_i \leftarrow f - t_e$ 
10      if  $f > t_e$  then
11        Replace the current occupant in cell  $e$  of the
        archive A with  $\theta_i$ 
12         $t_e \leftarrow (1 - \alpha)t_e + \alpha f$ 
13      end
14    end
15    rank  $\theta_i$  by  $\Delta_i$ 
16    Adapt CMA-ES parameters  $\theta, \Sigma, p$  based on
    improvement ranking  $\Delta_i$ 
17    if CMA-ES converges then
18      Restart CMA-ES with  $\Sigma = \sigma I$ .
19      Set  $\theta$  to a randomly selected existing cell  $\theta_i$  from
      the archive
20    end
21  end

```

Algorithm 1: Covariance Matrix Adaptation MAP-Annealing

The next theorem states that CMA-ME is equivalent to CMA-MAE when $\alpha = 1$ with the following caveats: First, we assume that CMA-ME restarts only by the CMA-ES restart rules, rather than the additional “no improvement” restart rule in prior work [26]. Second, we assume that both CMA-ME and CMA-MAE leverage μ selection [35] rather than filtering selection [26].

THEOREM 5.2. *The CMA-ME algorithm is equivalent to CMA-MAE when $\alpha = 1$ and $\min_f$ is an arbitrarily large negative number.*

We next provide theoretical insights on how the discount function f_A smoothly increases from a constant function $\min_f$ to the discount function used by CMA-ME, as α increases from 0 to 1. We focus on the special case of a fixed sequence of candidate solutions.

THEOREM 5.3. *Let α_i and α_j be two archive learning rates for archives A_i and A_j such that $0 \leq \alpha_i < \alpha_j \leq 1$. For two runs of CMA-MAE that generate the same sequence of m candidate solutions $\{S\} = \theta_1, \theta_2, \dots, \theta_m$, it follows that $f_{A_i}(\theta) \leq f_{A_j}(\theta)$ for all $\theta \in \mathbb{R}^n$.*

Finally, we wish to address the limitation that CMA-ME performs poorly on flat objectives, where all solutions have the same objective value. Consider how CMA-ME behaves on a flat objective $f(\theta) = C$ for all $\theta \in \mathbb{R}^n$, where C is an arbitrary constant. CMA-ME will only discover each new cell once and will not receive any further feedback from that cell, since any future solution cannot replace the incumbent elite. This hinders the CMA-ME’s movement in measure space, which is based on feedback from *changes* in the archive. Future candidate solutions will only fall into occupied cells, triggering repeated restarts caused by CMA-ES’s restart rule.

When the objective function plateaus, we still want CMA-ME to perform well and benefit from the CMA-ES adaptation mechanisms. One reasonable approach would be to keep track of the frequency o_e that each cell e has been visited in the archive, where o_e represents the number of times a solution was generated in that cell. Then, when a flat objective occurs, we rank solutions by descending frequency counts. Conceptually, CMA-ME would descend the density histogram defined by the archive, pushing the search towards regions of the measure space that have been sampled less frequently. Theorem 5.4 shows that we obtain the density descent behavior on flat objectives *without additional changes to the CMA-MAE algorithm*.

THEOREM 5.4. *The CMA-MAE algorithm optimizing a constant objective function $f(\theta) = C$ for all $\theta \in \mathbb{R}^n$ is equivalent to density descent, when $0 < \alpha < 1$ and $\min_f < C$.*

We highlight that the proof of Theorem 5.4 is based on two critical properties. First, the threshold update rule forms a strictly increasing sequence of thresholds for each cell. Second, CMA-ES is invariant to order preserving transformations of its objective f . While we have proposed the update rule of line 12 of Algorithm 1, we note that any update rule that satisfies the increasing sequence property retains the density descent property and is thus applicable in CMA-MAE.

While Theorem 5.4 assumes a constant objective f , we conjecture that the theorem holds true generally when threshold t_e in each cell e approaches the local optimum within the cell boundaries (see Conjecture E.7 in the Appendix).

6 EXPERIMENTS

We compare the performance of CMA-MAE with the state-of-the-art QD algorithms MAP-Elites, MAP-Elites (line), and CMA-ME, using existing Pyribs [77] QD library implementations. We set $\alpha = 0.01$ for CMA-MAE and include additional experiments for varying α in section 7. Because annealing methods replace solutions based on the threshold, we retain the best solution in each cell for comparison purposes. We include additional comparisons between CMA-MEGA and CMA-MAEGA – the DQD counterpart to CMA-MAE.

We select the benchmark domains from prior work [21]: linear projection [26], arm repertoire [16], and latent space illumination [24]. To evaluate the good exploration properties of CMA-MAE on flat objectives, we introduce a variant of the linear projection domain to include a “plateau” objective function that is constant everywhere for solutions within a fixed range and has a quadratic penalty for solutions outside the range. We describe the domains in detail in Appendix B.

We additionally introduce a second LSI experiment on StyleGAN2 [47], configured by insights from the generative art community [12, 27] that improve the quality of single-objective latent

	LP (sphere)		LP (Rastrigin)		LP (plateau)		Arm Repertoire		LSI (StyleGAN)		LSI (StyleGAN2)	
Algorithm	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage
MAP-Elites	41.64	50.80%	31.43	47.88%	47.07	47.07%	71.40	74.09%	12.85	19.42%	-936.96	4.48%
MAP-Elites (line)	49.07	60.42%	38.29	56.51%	52.20	52.20%	74.55	75.61%	14.40	21.11%	-236.65	8.81%
CMA-ME	36.50	42.82%	38.02	53.09%	34.54	34.54%	75.82	75.89%	14.00	19.57%	—	—
CMA-MAE	64.86	83.31%	52.65	80.46%	79.27	79.29%	79.03	79.24%	17.67	25.08%	—	—

Table 1: Mean QD-score and coverage values after 10,000 iterations for each QD algorithm per domain.

	LP (sphere)		LP (Rastrigin)		LP (plateau)		Arm Repertoire		LSI (StyleGAN)		LSI (StyleGAN2)	
Algorithm	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage
CMA-MEGA	75.32	100.00%	63.07	100.00%	100.00	100.00%	75.21	75.25%	16.08	22.58%	9.17	14.91%
CMA-MAEGA	75.39	100.00%	63.06	100.00%	100.00	100.00%	79.27	79.35%	16.20	23.83%	11.51	18.62%

Table 2: Mean QD-score and coverage values after 10,000 iterations for each DQD algorithm per domain.

space optimization. To improve control over image synthesis, the LSI (StyleGAN2) domain optimizes the full 9216-dimensional latent w -space, rather than a compressed 512-dimensional latent space in the LSI (StyleGAN) experiments. We exclude CMA-ME and CMA-MAE from this domain due to the prohibitive size of the covariance matrix. The LSI (StyleGAN2) domain allows us to evaluate the performance of DQD algorithms on a much more challenging DQD domain than prior work. We describe the domain in Appendix I.

6.1 Experiment Design

Independent Variables. We follow a between-groups design with two independent variables: the algorithm and the domain.

Dependent Variables. We use the sum of f values of all cells in the archive, defined as the QD-score [66], as a metric for the quality and diversity of solutions. Following prior work [21], we normalize the QD-score metric by the archive size (the total number of cells from the tessellation of measure space) to make the metric invariant to archive resolution. We additionally compute the coverage, defined as the number of occupied cells in the archive divided by the total number of cells.

6.2 Analysis

Derivative-free QD Algorithms. Table 1 shows the QD-score and coverage values for each algorithm and domain, averaged over 20 trials for the linear projection (LP) and arm repertoire domains and over 5 trials for the LSI domain. Fig. 3 shows the QD-score values for increasing number of iterations and example archives for CMA-MAE and CMA-ME with 95% confidence intervals.

We conducted a two-way ANOVA to examine the effect of the algorithm and domain on the QD-score. There was a significant interaction between the algorithm and the domain ($F(12, 320) = 1958.34, p < 0.001$). Simple main effects analysis with Bonferroni corrections showed that CMA-MAE outperformed all derivative-free QD baselines in all benchmark domains.

For the arm repertoire domain, we can compute the optimal archive coverage by testing whether each cell overlaps with a circle of radius equal to the maximum arm length (see Appendix B). We observe that CMA-MAE approaches the computed optimal coverage 80.24% for a resolution of 100×100 .

We observe from the runs that CMA-MAE initially explores regions of the measure space that have high-objective values. Once the archive becomes saturated, CMA-MAE reduces to approximate density descent, as we prove in Theorem 5.4 for flat objectives. On the other hand, CMA-ME does not receive any exploration signal when the objective landscape becomes flat, resulting in poor performance.

DQD Algorithms. We additionally compare CMA-MEGA and CMA-MAEGA in the five benchmark domains. Table 2 shows the QD-score and coverage values for each algorithm and domain, averaged over 20 trials for the linear projection (LP) and arm repertoire domains and over 5 trials for the LSI domains. We conducted a two-way ANOVA to examine the effect of the algorithm and domain on the QD-score. There was a significant interaction between the search algorithm and the domain ($F(5, 168) = 165.7, p < 0.001$). Simple main effects analysis with Bonferroni corrections showed that CMA-MAEGA outperformed CMA-MEGA in the LP (sphere), arm repertoire, and LSI (StyleGAN2) domains. There was no statistically significance difference between the two algorithms in the LP (Rastrigin), LP (plateau), and LSI (StyleGAN) domains.

We attribute the absence of a statistical difference in the QD-score between the two algorithms on the LP (Rastrigin) and LP (plateau) domains on the perfect coverage obtained by both algorithms. Thus, any differences in QD-score are based on the objective values of the solutions returned by each algorithm. In LP (plateau), the optimal objective for each cell is easily obtainable for both methods. The LP (Rastrigin) domain contains many local optima, because of the form of the objective function (Eq. 5). CMA-MEGA will converge to these optima before restarting, behaving as a single-objective optimizer within each local optimum. Because of the large number of local optima in the domain, CMA-MEGA obtains a higher QD-score.

In the LSI (StyleGAN) domain, we attribute similar performance between CMA-MEGA and CMA-MAEGA to the restart rules used to keep each search within the training distribution of StyleGAN. The ill-conditioned latent space of StyleGAN also explains why CMA-MAE outperforms both DQD algorithms on this domain. Being a natural gradient optimizer, CMA-MAE is an approximate second-order method, and second-order methods are better suited for optimizing spaces with ill-conditioned curvature.

On the other hand, in the LSI (StyleGAN2) domain, we regularize the search space by an L2 penalty in latent space, allowing for a




Figure 3: QD-score plot with 95% confidence intervals and heatmaps of generated archives by CMA-MAE and CMA-ME for the linear projection sphere (top), plateau (middle), and arm repertoire (bottom) domains. Each heatmap visualizes an archive of solutions across a 2D measure space.

larger learning rate and a basic restart rule for both algorithms, while still preventing drift out of the training distribution of StyleGAN2. Because of the fewer restarts, CMA-MAEGA can take advantage of the density descent property, which was shown to improve exploration in CMA-MAE, and outperform CMA-MEGA. Fig. 4 shows an

example collage generated from the final archive of CMA-MAEGA on the LSI (StyleGAN2) domain. We note that because StyleGAN2 has a better conditioning on the latent space [47], the model is better suited for first-order optimization of its latent space, which helps distinguish between the two QD algorithms.




Figure 4: A latent space illumination collage for the objective “A photo of the face of Tom Cruise.” with hair length and age measures sampled from a final CMA-MAEGA archive for the LSI (StyleGAN2) domain. See Appendix I for more detail.

We include runs for MAP-Elites and MAP-Elites (line) on the LSI (StyleGAN2) domain in Table 1 for comparison purposes. In the LSI (StyleGAN2) domain, the two algorithms drift out of distribution and suffer a regularization penalty that results in negative objective values. We observe that the gap in performance between derivative-free QD algorithms and QD algorithms is higher in the LSI (StyleGAN2) domain than in LSI (StyleGAN) domain. This highlights the benefits of leveraging gradients of the objective and measure functions in high-dimensional search spaces.

7 ON THE ROBUSTNESS OF CMA-MAE

Next, we present two studies that evaluate CMA-MAE robustness across two hyperparameters that may affect algorithm performance: the archive learning rate α and the archive resolution.

Archive Learning Rate. We examine the effect of different archive learning rates on the performance of CMA-MAE in the linear projection and arm repertoire domains. We vary the learning rate from 0 to 1 on an exponential scale, while keeping the resolution constant.

Table 3 shows that running CMA-MAE with the different $0 < \alpha < 1$ values results in similar performance, showing that CMA-MAE is fairly robust to the exact choice of α value. On the other hand, if $\alpha = 0$ or $\alpha = 1$ the performance drops drastically. Setting $\alpha = 1$ results in very similar performance with CMA-ME, which supports our insight from Theorem 5.2.

Archive Resolution. As noted in prior work [14, 21], quality diversity algorithms in the MAP-Elites family sometimes perform differently when run with different archive resolutions. For example, in the linear projection domain proposed in prior work [26], CMA-ME outperformed MAP-Elites and MAP-Elites (line) for archives of resolution 500×500 , while in this paper we observe that it performs worse for resolution 100×100 . In this study, we investigate how CMA-MAE performs at different archive resolutions.

First, we note that the optimal archive learning rate α is dependent on the resolution of the archive. Consider as an example a sequence of solution additions to two archives A_1 and A_2 of resolution 100×100 and 200×200 , respectively. A_2 subdivides each cell in A_1 into four cells, thus archive A_2 ’s thresholds t_e should increase at a four times faster rate than A_1 . To account for this difference, we compute α_2 for A_2 via a conversion formula $\alpha_2 = 1 - (1 - \alpha_1)^r$ (see derivation in Appendix G), where r is the ratio of cell counts between archives A_1 and A_2 . We initialize $\alpha_1 = 0.01$ for A_1 . In the above example, $\alpha_2 = 1 - (1 - 0.01)^4 = 0.0394$.

Fig. 5 shows the QD-score of CMA-MAE with the resolution-dependent archive learning rate and the baselines for each benchmark domain. CMA-ME performs worse as the resolution decreases because the archive changes quickly at small resolutions, affecting CMA-ME’s adaptation mechanism. On the contrary, MAP-Elites and MAP-Elites (line) perform worse as the resolution increases




Figure 5: Final QD-score of each algorithm for 25 different archive resolutions.

	LP (sphere)		LP (Rastrigin)		LP (plateau)		Arm Repertoire	
α (CMA-MAE)	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage	QD-score	Coverage
0.000	5.82	6.06%	5.33	6.24%	19.49	19.49%	65.91	66.25%
0.001	62.65	79.36%	47.87	68.10%	77.60	77.68%	78.63	79.07%
0.010	64.86	83.31%	52.65	80.56%	79.27	79.29%	79.03	79.24%
0.100	60.42	76.19%	48.74	72.50%	83.21	83.21%	78.74	78.85%
1.000	37.01	43.50%	37.86	52.82%	34.00	34.00%	75.94	76.01%

Table 3: Mean QD metrics after 10,000 iterations for CMA-MAE at different learning rates.

due to having more elites to perturb. CMA-MAE’s performance is invariant to the resolution of the archive.

8 RELATED WORK

Quality Diversity Optimization. The predecessor to quality diversity optimization, simply called diversity optimization, originated with the Novelty Search algorithm [54], which searches for a collection of solutions that are diverse in measure space. Later work introduced the Novelty Search with Local Competition (NSLC) [55] and MAP-Elites [15, 59] algorithms, which combined single-objective optimization with diversity optimization and were the first QD algorithms. Since then, several QD algorithms have been proposed, based on a variety of single-objective optimization methods, such as Bayesian optimization [49], evolution strategies [10, 11, 26], differential evolution [9], and gradient ascent [21]. Several works have improved selection mechanisms [16, 71], archives [23, 72, 78], perturbation operators [61, 79], and resolution scaling [13, 23, 32].

QD with Gradient Information. Several works combine gradient information with QD optimization without leveraging the objective and measure gradients directly. For example, in model-based QD optimization [7, 28, 29, 34, 48, 57, 80], prior work [68] trains an autoencoder on the archive of solutions and leverages the Jacobian of the decoder network to compute the covariance of the Gaussian perturbation. Works in quality diversity reinforcement learning (QD-RL) [60, 63, 65, 76] approximate a reward gradient or diversity gradient via a critic network, action space noise, or evolution strategies and incorporate those gradients into a QD-RL algorithm.

Acceptance Thresholds. Our proposed archive learning rate α was loosely inspired by simulated annealing methods [2] that maintain an acceptance threshold that gradually becomes more selective as the algorithm progresses. The notion of an acceptance threshold is also closely related to minimal criterion methods in evolutionary computation [5, 6, 53, 73]. Our work differs by both 1) maintaining an acceptance threshold per archive cell rather than a global threshold and 2) annealing the threshold.

9 LIMITATIONS AND FUTURE WORK

Our approach introduced two hyperparameters, α and $\min_f$ (see Appendix J for $\min_f$ analysis), to control the rate that $f - f_A$ changes. We observed that an α set strictly between 0 and 1 yields theoretical exploration improvements and that CMA-MAE is robust with respect to the exact choice of α . We additionally derived a conversion formula that converts an α_1 for a specific archive resolution to an equivalent α_2 for a different resolution. However, the conversion formula still requires practitioners to specify a good initial value of α_1 . Future work will explore ways to automatically initialize α , similar to how CMA-ES automatically assigns internal parameters [35].

CMA-MAE’s DQD counterpart CMA-MAEGA sets a new state-of-the-art in DQD optimization. However, observing its performance benefits required the more challenging LSI (StyleGAN2) domain. This highlights the need for more challenging DQD problems to advance research in DQD algorithms, since many of the current benchmark domains can be solved optimally by existing algorithms.

Quality diversity optimization is a rapidly growing branch of stochastic optimization with applications in generative design [28, 29, 33], automatic scenario generation in robotics [20, 22, 25], reinforcement learning [60, 63, 65, 76], damage recovery in robotics [15], and procedural content generation [4, 17, 24, 31, 50, 69, 70, 75, 80]. Our paper introduces a new quality diversity algorithm, CMA-MAE. Our theoretical findings inform our experiments, which show that CMA-MAE addresses three major limitations affecting CMA-ME, leading to state-of-the-art performance and robustness.

10 ACKNOWLEDGMENTS

This work was supported by the NSF CAREER Award (#2145077) and NSF NRI (#2024949). We thank Bryon Tjanaka and David Lee for their helpful discussions and feedback when deriving the batch threshold update rule in Appendix H. We would also like to thank Lisa B. Soros for her feedback on a preliminary version of this paper.

REFERENCES

- [1] AKIMOTO, Y., NAGATA, Y., ONO, I., AND KOBAYASHI, S. Bidirectional relation between cma evolution strategies and natural evolution strategies. In *International Conference on Parallel Problem Solving from Nature* (2010), Springer, pp. 154–163.
- [2] BERTSIMAS, D., AND TSITSIKLIS, J. Simulated annealing. *Statistical science* 8, 1 (1993), 10–15.
- [3] BESNIER, V., JAIN, H., BURSUC, A., CORD, M., AND PÉREZ, P. This dataset does not exist: training models from generated images. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020), IEEE, pp. 1–5.
- [4] BHATT, V., TJANAKA, B., FONTAINE, M., AND NIKOLAIDIS, S. Deep surrogate assisted generation of environments. *Advances in Neural Information Processing Systems* (2022).
- [5] BRANT, J. C., AND STANLEY, K. O. Minimal criterion coevolution: a new approach to open-ended search. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2017), pp. 67–74.
- [6] BRANT, J. C., AND STANLEY, K. O. Diversity preservation in minimal criterion coevolution through resource limitation. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (2020), pp. 58–66.
- [7] CAZENILLE, L., BREDECHE, N., AND AUBERT-KATO, N. Exploring self-assembling behaviors in a swarm of bio-micro-robots using surrogate-assisted map-elites. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)* (2019), IEEE, pp. 238–246.
- [8] CHATZILYGEROUDIS, K., CULLY, A., VASSILIADES, V., AND MOURET, J.-B. Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer, 2021, pp. 109–135.
- [9] CHOI, T. J., AND TOGELIUS, J. *Self-Referential Quality Diversity through Differential MAP-Elites*. Association for Computing Machinery, New York, NY, USA, 2021, p. 502–509.
- [10] COLAS, C., MADHAVAN, V., HUIZINGA, J., AND CLUNE, J. Scaling map-elites to deep neuroevolution. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (2020), pp. 67–75.
- [11] CONTI, E., MADHAVAN, V., PETROSKI SUCH, F., LEHMAN, J., STANLEY, K., AND CLUNE, J. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *Advances in neural information processing systems* 31 (2018).
- [12] CROWSON, K., BIDERMAN, S., KORNIS, D., STANDER, D., HALLAHAN, E., CASTRICATO, L., AND RAFF, E. Vqgan-clip: Open domain image generation and editing with natural language guidance. *arXiv preprint arXiv:2204.08583* (2022).
- [13] CULLY, A. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2019), pp. 81–89.
- [14] CULLY, A. Multi-emitter map-elites: improving quality, diversity and data efficiency with heterogeneous sets of emitters. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2021), pp. 84–92.
- [15] CULLY, A., CLUNE, J., TARAPORE, D., AND MOURET, J.-B. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503.
- [16] CULLY, A., AND DEMIRIS, Y. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation* 22, 2 (2017), 245–259.
- [17] EARLE, S., SNIDER, J., FONTAINE, M. C., NIKOLAIDIS, S., AND TOGELIUS, J. Illuminating diverse neural cellular automata for level generation. *arXiv preprint arXiv:2109.05489* (2021).
- [18] ECOFFET, A., CLUNE, J., AND LEHMAN, J. Open questions in creating safe open-ended ai: tensions between control and creativity. *arXiv preprint arXiv:2006.07495* (2020).
- [19] FEFFERMAN, C., MITTER, S., AND NARAYANAN, H. Testing the manifold hypothesis. *Journal of the American Mathematical Society* 29, 4 (2016), 983–1049.
- [20] FONTAINE, M., HSU, S., ZHANG, Y., TJANAKA, B., AND NIKOLAIDIS, S. On the importance of environments in human-robot coordination. *Robotics: Science and Systems (RSS)* (2021).
- [21] FONTAINE, M., AND NIKOLAIDIS, S. Differentiable quality diversity. *Advances in Neural Information Processing Systems* 34 (2021).
- [22] FONTAINE, M., AND NIKOLAIDIS, S. A quality diversity approach to automatically generating human-robot interaction scenarios in shared autonomy. *Robotics: Science and Systems* (2021).
- [23] FONTAINE, M. C., LEE, S., SOROS, L. B., DE MESENTIER SILVA, F., TOGELIUS, J., AND HOOVER, A. K. Mapping hearthstone deck spaces through map-elites with sliding boundaries. In *Proceedings of The Genetic and Evolutionary Computation Conference* (2019), pp. 161–169.
- [24] FONTAINE, M. C., LIU, R., TOGELIUS, J., HOOVER, A. K., AND NIKOLAIDIS, S. Illuminating mario scenes in the latent space of a generative adversarial network. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2021).
- [25] FONTAINE, M. C., AND NIKOLAIDIS, S. Evaluating human-robot interaction algorithms in shared autonomy via quality diversity scenario generation. *International Journal of Human-Robot Interaction* (jul 2021).
- [26] FONTAINE, M. C., TOGELIUS, J., NIKOLAIDIS, S., AND HOOVER, A. K. Covariance matrix adaptation for the rapid illumination of behavior space. In *Proceedings of the 2020 genetic and evolutionary computation conference* (2020), pp. 94–102.
- [27] FRANS, K., SOROS, L., AND WITKOWSKI, O. Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. *arXiv preprint arXiv:2106.14843* (2021).
- [28] GAIER, A., ASTEROTH, A., AND MOURET, J.-B. Data-efficient design exploration through surrogate-assisted illumination. *Evolutionary computation* 26, 3 (2018), 381–410.
- [29] GAIER, A., ASTEROTH, A., AND MOURET, J.-B. Discovering representations for black-box optimization. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2020), GECCO '20, Association for Computing Machinery, p. 103–111.
- [30] GRAHAM, R. L., KNUTH, D. E., PATASHNIK, O., AND LIU, S. Concrete mathematics: a foundation for computer science. *Computers in Physics* 3, 5 (1989), 106–107.
- [31] GRAVINA, D., KHALIFA, A., LIAPIS, A., TOGELIUS, J., AND YANNAKAKIS, G. N. Procedural content generation through quality diversity. In *2019 IEEE Conference on Games (CoG)* (2019), IEEE, pp. 1–8.
- [32] GRILLOTTI, L., AND CULLY, A. Relevance-guided unsupervised discovery of abilities with quality-diversity algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2022), pp. 77–85.
- [33] HAGG, A., BERNIS, S., ASTEROTH, A., COLTON, S., AND BÄCK, T. Expressivity of parameterized and data-driven representations in quality diversity search. In *Proceedings of the Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2021), GECCO '21, Association for Computing Machinery, p. 678–686.
- [34] HAGG, A., WILDE, D., ASTEROTH, A., AND BÄCK, T. Designing air flow with surrogate-assisted phenotypic niching. In *International Conference on Parallel Problem Solving from Nature* (2020), Springer, pp. 140–153.
- [35] HANSEN, N. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).
- [36] HANSEN, N., AUGER, A., MERSMANN, O., TUSAR, T., AND BROCKHOFF, D. Coco: A platform for comparing continuous optimizers in a black-box setting.
- [37] HANSEN, N., AUGER, A., ROS, R., FINCK, S., AND POŠÍK, P. Comparing results of 31 algorithms from the black-box optimization benchmarking bbo-2009. pp. 1689–1696.
- [38] HANSEN, N., MÜLLER, S. D., AND KOUMOUTSAKOS, P. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary computation* 11, 1 (2003), 1–18.
- [39] HANSEN, N., AND OSTERMEIER, A. Convergence properties of evolution strategies with the derandomized covariance matrix adaptation: The (μ)-es. *Eufit* 97 (1997), 650–654.
- [40] HENDRYCKS, D., CARLINI, N., SCHULMAN, J., AND STEINHARDT, J. Unsolved problems in ml safety. *arXiv preprint arXiv:2109.13916* (2021).
- [41] JAHANIAN, A., PUIG, X., TIAN, Y., AND ISOLA, P. Generative models as a data source for multiview representation learning. *arXiv preprint arXiv:2106.05258* (2021).
- [42] JAIN, N., OLMO, A., SENGUPTA, S., MANIKONDA, L., AND KAMBHAMPATI, S. Imperfect imagination: Implications of gans exacerbating biases on facial data augmentation and snapchat selfie lenses. *arXiv preprint arXiv:2001.09528* (2020).
- [43] JOHNSON, N. L., KOTZ, S., AND BALAKRISHNAN, N. *Continuous univariate distributions, volume 2*, vol. 289. John Wiley & sons, 1995.
- [44] KARRAS, T., AITTALA, M., HELLSTEN, J., LAINE, S., LEHTINEN, J., AND AILA, T. Training generative adversarial networks with limited data. In *Proc. NeurIPS* (2020).
- [45] KARRAS, T., AITTALA, M., LAINE, S., HÄRKÖNEN, E., HELLSTEN, J., LEHTINEN, J., AND AILA, T. Alias-free generative adversarial networks. In *Proc. NeurIPS* (2021).
- [46] KARRAS, T., LAINE, S., AND AILA, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4401–4410.
- [47] KARRAS, T., LAINE, S., AITTALA, M., HELLSTEN, J., LEHTINEN, J., AND AILA, T. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020).
- [48] KELLER, L., TANNEBERG, D., STARK, S., AND PETERS, J. Model-based quality-diversity search for efficient robot learning. *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2020), 9675–9680.
- [49] KENT, P., AND BRANKE, J. Bop-elites, a bayesian optimisation algorithm for quality-diversity search. *arXiv preprint arXiv:2005.04320* (2020).
- [50] KHALIFA, A., LEE, S., NEALEN, A., AND TOGELIUS, J. Talakat: Bullet hell generation through constrained map-elites. In *Proceedings of The Genetic and Evolutionary Computation Conference* (2018), pp. 1047–1054.
- [51] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. In *3rd*

- International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015).
- [52] LEHMAN, J., CLUNE, J., MISEVIC, D., ADAMI, C., ALTENBERG, L., BEAULIEU, J., BENTLEY, P. J., BERNARD, S., BESLON, G., BRYSON, D. M., ET AL. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life* 26, 2 (2020), 274–306.
 - [53] LEHMAN, J., AND STANLEY, K. O. Revising the evolutionary computation abstraction: minimal criteria novelty search. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation* (2010), pp. 103–110.
 - [54] LEHMAN, J., AND STANLEY, K. O. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011), 189–223.
 - [55] LEHMAN, J., AND STANLEY, K. O. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation* (2011), pp. 211–218.
 - [56] LIM, B., ALLARD, M., GRILLOTTI, L., AND CULLY, A. Accelerated quality-diversity for robotics through massive parallelism. In *ICLR Workshop on Agent Learning in Open-Endedness* (2022).
 - [57] LIM, B., GRILLOTTI, L., BERNASCONI, L., AND CULLY, A. Dynamics-aware quality-diversity for efficient learning of skill repertoires. *ArXiv abs/2109.08522* (2021).
 - [58] MENON, S., DAMIAN, A., HU, S., RAVI, N., AND RUDIN, C. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 2437–2445.
 - [59] MOURET, J.-B., AND CLUNE, J. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
 - [60] NILSSON, O., AND CULLY, A. Policy gradient assisted map-elites. *Proceedings of the Genetic and Evolutionary Computation Conference* (2021).
 - [61] NORDMOEN, J., SAMUELSEN, E., ELLEFSEN, K. O., AND GLETTE, K. Dynamic mutation in map-elites for robotic repertoire generation. In *Artificial Life Conference Proceedings* (2018), MIT Press, pp. 598–605.
 - [62] PAOLO, G., CONINX, A., DONCIEUX, S., AND LAFLAQUIÈRE, A. Sparse reward exploration via novelty search and emitters. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2021), pp. 154–162.
 - [63] PARKER-HOLDER, J., PACCHIANO, A., CHOROMANSKI, K. M., AND ROBERTS, S. J. Effective diversity in population based reinforcement learning. In *Advances in Neural Information Processing Systems* (2020), H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33, Curran Associates, Inc., pp. 18050–18062.
 - [64] PEREZ, V. Generating images from prompts using clip and stylegan, 2021. MIT License.
 - [65] PIERROT, T., MACÉ, V., CIDERON, G., BEGUIR, K., CULLY, A., SIGAUD, O., AND PERRIN-GILBERT, N. Diversity policy gradient for sample efficient quality-diversity optimization. *arXiv e-prints* (2020), arXiv:2006.
 - [66] PUGH, J. K., SOROS, L. B., SZERLIP, P. A., AND STANLEY, K. O. Confronting the challenge of quality diversity. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation* (2015), pp. 967–974.
 - [67] RADFORD, A., KIM, J. W., HALLACY, C., RAMESH, A., GOH, G., AGARWAL, S., SASTRY, G., ASKELL, A., MISHKIN, P., CLARK, J., KRUEGER, G., AND SUTSKEVER, I. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning ICML 2021* (2021).
 - [68] RAKICEVIC, N., CULLY, A., AND KORMUSHEV, P. Policy manifold search: Exploring the manifold hypothesis for diversity-based neuroevolution. In *Proceedings of the Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2021), GECCO '21, Association for Computing Machinery, p. 901–909.
 - [69] SARKAR, A., AND COOPER, S. Generating and blending game levels via quality-diversity in the latent space of a variational autoencoder. *arXiv preprint arXiv:2102.12463* (2021).
 - [70] SCHRUM, J., VOLZ, V., AND RISI, S. Cppn2gan: Combining compositional pattern producing networks and gans for large-scale pattern generation. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (2020), pp. 139–147.
 - [71] SFIKAS, K., LIAPIS, A., AND YANNAKAKIS, G. N. Monte carlo elites: Quality-diversity selection as a multi-armed bandit problem. In *Proceedings of the Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2021), GECCO '21, Association for Computing Machinery, p. 180–188.
 - [72] SMITH, D., TOKARCHUK, L., AND WIGGINS, G. Rapid phenotypic landscape exploration through hierarchical spatial partitioning. In *International conference on parallel problem solving from nature* (2016), Springer, pp. 911–920.
 - [73] STANLEY, K. O., CHENEY, N., AND SOROS, L. How the strictness of the minimal criterion impacts open-ended evolution. In *ALIFE 2016, the Fifteenth International Conference on the Synthesis and Simulation of Living Systems* (2016), MIT Press, pp. 208–215.
 - [74] STANLEY, K. O., LEHMAN, J., AND SOROS, L. Open-endedness: The last grand challenge you’ve never heard of. *While open-endedness could be a force for discovering intelligence, it could also be a component of AI itself* (2017).
 - [75] STECKEL, K., AND SCHRUM, J. Illuminating the space of beatable lode runner levels produced by various generative adversarial networks. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (New York, NY, USA, 2021), GECCO '21, Association for Computing Machinery, p. 111–112.
 - [76] TJANAKA, B., FONTAINE, M. C., TOGELIUS, J., AND NIKOLAIDIS, S. Approximating gradients for differentiable quality diversity in reinforcement learning. *GECCO* (2022).
 - [77] TJANAKA, B., FONTAINE, M. C., ZHANG, Y., SOMMERER, S., DENNLER, N., AND NIKOLAIDIS, S. pyribs: A bare-bones python library for quality diversity optimization. <https://github.com/icaros-usc/pyribs>, 2021.
 - [78] VASSILIADES, V., CHATZILYGEROUDIS, K., AND MOURET, J.-B. Using centroidal voronoi tessellations to scale up the multidimensional archive of phenotypic elites algorithm. *IEEE Transactions on Evolutionary Computation* 22, 4 (2018), 623–630.
 - [79] VASSILIADES, V., AND MOURET, J.-B. Discovering the elite hypervolume by leveraging interspecies correlation. In *Proceedings of the Genetic and Evolutionary Computation Conference* (2018), pp. 149–156.
 - [80] ZHANG, Y., FONTAINE, M. C., HOOVER, A. K., AND NIKOLAIDIS, S. Deep surrogate assisted map-elites for automated hearthstone deckbuilding. *arXiv preprint arXiv:2112.03534* (2021).