

GRADIENT-INFORMED QUALITY DIVERSITY FOR THE ILLUMINATION OF DISCRETE SPACES

Anonymous authors

Paper under double-blind review

ABSTRACT

Quality Diversity (QD) algorithms have been proposed to search for a large collection of both diverse and high-performing solutions instead of a single set of local optima. While early QD algorithms view the objective and descriptor functions as black-box functions, novel tools have been introduced to use gradient information to accelerate the search and improve overall performance of those algorithms over continuous input spaces. However a broad range of applications involve discrete spaces, such as drug discovery or image generation. Exploring those spaces is challenging as they are combinatorially large and gradients cannot be used in the same manner as in continuous spaces. We introduce MAP-ELITES with a Gradient-Informed Discrete Emitter (ME-GIDE), which extends QD optimisation with differentiable functions over discrete search spaces. ME-GIDE leverages the gradient information of the objective and descriptor functions with respect to its discrete inputs to propose gradient-informed updates that guide the search towards a diverse set of high quality solutions. We evaluate our method on challenging benchmarks including protein design and discrete latent space illumination and find that our method outperforms state-of-the-art QD algorithms in all benchmarks.

1 INTRODUCTION

Quality-Diversity (QD) Optimization algorithms Chatzilygeroudis et al. (2021); Cully & Demiris (2017) have changed the classical paradigm of optimization: inspired by natural evolution, the essence of QD methods is to provide a large and diverse set of high-performing solutions rather than only the best one. This core idea showed great outcomes in different fields such as robotics Cully et al. (2015) where it allows to learn to control robots using diverse and efficient policies or latent space exploration Fontaine et al. (2021) of generative models to generate a diverse set of high quality images or video game levels. The main QD approaches are derivative-free optimizers, such as the MAP-ELITES algorithm Mouret & Clune (2015), a genetic algorithm that tries to find high quality solutions covering the space defined by a variety of user-defined features of interest. However, subsequent methods Nilsson & Cully (2021); Fontaine & Nikolaidis (2021) have shown that when the objective functions (quality) or the descriptor functions (diversity) are differentiable, using gradients of those functions can improve convergence both in terms of speed and performance compared to traditional QD algorithms. Those methods were applied to Reinforcement Learning problems such as robotics Pierrot et al. (2022a) or latent space illumination Fontaine & Nikolaidis (2021) of generative adversarial networks.

Those applications focused on optimization over continuous variables. However many real-world applications are best framed using discrete features. For instance, generative architecture such as discrete VAEs Van Den Oord et al. (2017); Rolfe (2016) have shown capability to generate high-quality images Razavi et al. (2019) and have been at the core of recent successful approaches for text-to-image generation Ramesh et al. (2021); Gafni et al. (2022). A further natural application using discrete variables is protein design Huang et al. (2016): proteins play a role in key functions in nature and protein design allows to create new drugs or biofuels. As proteins are sequences of 20 possible amino acids, designing them is a tremendously hard task as only a few of those possible sequences are plausible in the real world. Deep learning for biological sequences provided significant advances in several essential tasks for biological sequence, such AlphaFold for structure prediction Jumper et al. (2021a) or language models for amino acid likelihood prediction Rives et al.

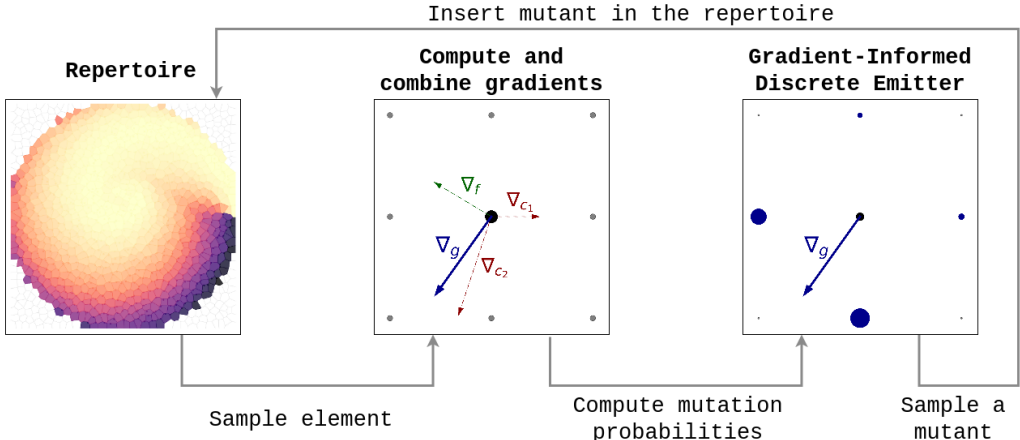


Figure 1: MAP-ELITES with Gradients Informed Discrete Emitter (ME-GIDE). At each iteration, an element is sampled in the repertoire. Gradients are computed over continuous fitness and descriptor functions with respect to their discrete inputs. Gradients are linearly combined to favour higher fitness and exploration of the descriptor space. Probabilities of mutation over the neighbours of the element are derived from this gradient information. Finally, a mutant is sampled according to those probabilities and inserted back in the repertoire.

(2021) . These successes motivate the use of these tools as objective functions for protein design Anishchenko et al. (2021a).

In practice, objective functions are often written as differentiable functions of their inputs, enabling the use of gradient information of the continuous extension if the input are discrete Grathwohl et al. (2021). We second this research trend and address gradient informed quality-diversity optimization over discrete inputs when objective and descriptors functions are differentiable. We make several contributions: (i) we introduce MAP-ELITES with a Gradient-Informed Discrete Emitter (ME-GIDE), a genetic algorithm where mutations are sampled thanks to a gradient informed distribution. (ii) We propose a way to select the hyperparameters of this algorithm by weighting the importance we give to this gradient information. (iii) We demonstrate the ability of our method to better illuminate discrete spaces on proteins and the discrete latent space of a generative model.

2 PROBLEM DEFINITION

In this work, we consider the problem of Quality Diversity Optimization over discrete spaces. The QD problem assumes an objective function $f : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is called search space, and d descriptors $c_i : \mathcal{X} \rightarrow \mathbb{R}$, or as a single descriptor function $\mathbf{c} : \mathcal{X} \rightarrow \mathbb{R}^d$. We note $S = \mathbf{c}(\mathcal{X})$ the descriptor space formed by the range of \mathbf{c} . We only consider discrete spaces of dimension m with K categories such that $\mathcal{X} \subset \{1, \dots, K\}^m$. QD algorithms of the family of the MAP-ELITES algorithm, discretize the descriptor space S via a tessellation method. Let \mathcal{T} be the tessellation of S into M cells S_i . The goal of QD methods is to find a set of solutions $\mathbf{x}_i \in \mathcal{X}$ so that each solution \mathbf{x}_i occupies a different cell S_i in \mathcal{T} and maximizes the objective function within that cell. The QD objective can thus be formalized as follows:

$$\max_{\mathbf{x} \in \mathcal{X}} \sum_{i=1}^M f(\mathbf{x}_i), \text{ where } \forall i, \mathbf{c}(\mathbf{x}_i) \in S_i. \tag{1}$$

We consider here that both the objective and descriptor functions are actually defined on real values \mathbb{R} and restricted to the discrete inputs \mathcal{X} . We also consider them to be first-order differentiable, hence for any input $x \in \mathcal{X}$, we can compute gradients $\nabla_x f(x)$ and $\nabla_x c_i(x)$.

Algorithm 1: Gradient Informed Discrete Emitter (GIDE)

Given: a batch of B elements x_n from the repertoire.
for $1 \leq n \leq B$ **do**
 // Compute the distribution
 Compute $\tilde{\mathbf{d}}_n$ using Equation 2.
 Compute the proposal distribution p_n using Equation 4.

 // Target entropy
 Adjust T to obtain a target entropy $\bar{H} = \bar{H}_{\text{target}}$ over the proposal distributions.

 // Sample
 Draw a mutant x'_n from each x_n according to p_n .
end
Return x'_1, \dots, x'_B

3 MAP-ELITES WITH GRADIENT PROPOSAL EMITTER

MAP-Elites and Differentiable MAP-Elites MAP-ELITES is a QD method that discretizes the space of possible descriptors into a repertoire (also called archive) via a tessellation method. Its goal is to fill each cell of this repertoire with the highest performing individuals. To do so, MAP-ELITES firstly initializes a repertoire over the BD space and secondly initializes random solutions, evaluates them and inserts them in the repertoire. Then successive iterations are performed: (i) select and copy solutions uniformly over the repertoire (ii) mutate the copies to create new solution candidates (iii) evaluate the candidates to determine their descriptor and objective value (iv) find the corresponding cells in the tessellation (v) if the cell is empty, introduce the candidate and otherwise replace the solution already in the cell by the new candidate if it has a greater objective function value. When using real variables, the most popular choice is the Iso+LineDD Vassiliades & Mouret (2018) that mixes two gaussian perturbations. When using discrete variables, mutations are generally defined as point mutation, *ie* select a random position and flip its value from the current one to a new one.

To incorporate gradients into MAP-ELITES algorithms, (Fontaine & Nikolaidis, 2021) introduced MAP-ELITES with a Gradient Arborecence (MEGA) a novel way to use gradient information to guide the mutations. First, the authors propose a novel objective function to encompass both quality and diversity: $g(x) = f(x) + \sum_{i=1}^d w_i c_i(x)$, where $w_i \sim \mathcal{N}(0, \sigma_g I) \forall i \in \{0, \dots, d\}$. In one version of their algorithm, OMG-MEGA, authors simply extend MAP-ELITES by mutating selected element of the repertoires x via $x' = x + |w_0| \nabla_x g(x)$ where w_i are sampled at each iteration and for each element in the batch and σ_g acts similarly to a learning rate. Indeed, maximizing g will direct the mutations towards higher fitness and different directions of the descriptor space thanks to the randomness introduced by sampling the w_i .

From Gradients to Gradient Informed Discrete Emitter As this approach has proven effective on several tasks, we follow the previous formulation by trying to maximize g . In this case at a given iteration, for a given sampled $x \in \mathcal{X}$ and given sampled coefficients c_0, \dots, c_d , our mutation should ideally find a neighbour $x' \in \mathcal{X}$ that maximizes g . Ideally, our operator should provide $x' = \arg \max_{z \in H_\tau(x)} g(z) - g(x)$. where $H_\tau(x)$ the Hamming ball of size τ around x our mutation.

However, for a given variable $x \in \{1, \dots, K\}^m$, even for $\tau = 1$ the cardinality of this Hamming ball is $mK - 1$. Hence finding the optimal x' requires $mK - 1$ evaluations of g . Doing it at each step would be too expensive but since those differences are local, they can be approximated using gradient information. Following Grathwohl et al. (2021) we use Taylor-series approximation of the local differences: noting $x^{(i,k)}$ the mutant from x by flipping the position i from x_i to k , we estimate those differences as

$$g(x^{(i,k)}) - g(x) \simeq \nabla_x g(x)_{ik} - x_i^T \nabla_x g(x)_i = \tilde{d}_{ik} \quad (2)$$

where x' differs from x by a flip on position $1 \leq i \leq m$ from x_i to $k \in \{1, \dots, K\}$. The previous formulation works in the case where $\tau = 1$ but similar approximations can be derived for larger window sizes at the expense of the approximation's quality.

Using those approximations, we could straightforwardly compute a local maximum of g by taking the argmax at each step. In order to encourage exploration, we use those approximate local differences to create a proposal distribution over the neighbours of x . Inspired by sampling techniques such as Metropolis-Hastings, we define the probability $p(x^{(i,k)}|x)$ of mutating from x to $x^{(i,k)}$ as

$$p(x^{(i,k)}|x) \propto e^{\frac{\tilde{d}_{ik}}{T}} \quad (3)$$

where $T > 0$ is a temperature parameter. The temperature parameter tempers whether the mutation will go towards the argmax ($T \rightarrow 0$) or a purely random mutation ($T \rightarrow +\infty$). Finally, in order to sample flips at position i from x_i to k , we normalize those probabilities by computing a softmax over every possible flip (mK possibilities):

$$p(x^{(i,k)}|x) = \frac{e^{\tilde{d}_{ik}/T}}{\sum_{i,l=1}^{m,K} e^{\tilde{d}_{jl}/T}} \quad (4)$$

As every approximated difference \tilde{d}_{ik} can be computed in one gradient evaluation $\nabla_x g(x)$, those probabilities can be efficiently computed.

Formally our Gradient Informed Discrete Emitter (GIDE) receives a candidate x , computes \tilde{d}_{ik} using Equation 2, then computes mutation probabilities using Equation 4 and finally samples a mutated $x^{(i,k)}$ using these probabilities. We summarize this procedure in Algorithm 1.

Using this emitter, we design our main algorithm: MAP-ELITES with a Gradient Informed Discrete Emitter (ME-GIDE). The procedure follows the one of MAP-ELITES: we firstly initialize a repertoire of behaviour descriptors. Then at each iteration, the following operations are performed: (i) select a batch of solutions uniformly over the repertoire (ii) sample mutants using our GIDE (iii) evaluate the candidates and add them in the repertoire if their fitness is higher than the existing solution in the cell. This procedure is described in Algorithm 2 with an additional step to control the strength of the gradient guidance defined in the following paragraph.

Algorithm 2: MAP-Elites with Gradient Informed Discrete Emitter (ME-GIDE)

Given: the number of cells M and tessellation \mathcal{T} ; the batch size B and the number of iterations N ; the descriptors function \mathbf{c} and the multi-objective function \mathbf{f} ; a target entropy value H_{target} ; the initial population of solution candidates $\{\mathbf{x}_k\}$

```
// Initialization
```

```
For each initial solution, find the cell corresponding to its descriptor and add initial solutions to their cells.
```

```
// Main loop
```

```
for  $1 \leq n_{steps} \leq N$  do
```

```
    // Select new generation
```

```
    Sample uniformly  $B$  solutions  $x_n$  in the repertoire with replacement
```

```
    // Compute the gradients
```

```
    Randomly draw weights  $\mathbf{w}^{(n)} \sim \mathcal{N}(0, I)$ 
```

```
    Compute gradients  $\nabla_{x_n} f(x_n)$  and  $\nabla_{x_n} c(x_n)$ 
```

```
    Normalize the gradients.
```

```
    Compute combined updates  $\nabla_n = |w_0^{(n)}| \nabla_{x_n} f(x_n) + \sum_{i=1}^n w_i^{(n)} \nabla_{x_n} c_i(x_n)$ 
```

```
    // Gradient-Informed Discrete Emitter
```

```
    Use Algorithm 1 to sample mutants  $x'_n$ 
```

```
    // Addition in the archive
```

```
    Add each mutant in its corresponding cell if it improves the cell fitness, otherwise discard it.
```

```
end
```

Controlling the Gradient Guidance with a Target Entropy The shape of the flip distribution $p_{i,k}$ computed by GIDE strongly depends on the choice of the temperature parameter T . However, searching for an optimal value of this unbound parameter can be tedious owing to the fact that it lies in $[0, +\infty]$ and has limited human-interpretability. Since this temperature parameter T controls the concentration of the probability mass over the directions with highest estimated-improvement values $\tilde{d}_{i,k}$, mispecifying it would result in suffering either degenerated or uniform distribution. The former means being too greedy towards the most promising directions, which hinders exploration, while the latter means under-exploiting our estimate. We propose to tackle this exploration-exploitation tradeoff by controlling the Shannon entropy of the flip-distribution $H(p_{i,k})$. When $T \rightarrow 0$, we approach the deterministic distribution over the highest estimate, i.e. with minimum entropy and when $T \rightarrow +\infty$ we approach the uniform distribution, i.e. with maximum entropy.

To avoid either degenerate cases, we set a target \bar{H}_{target} for the normalized entropy $\bar{H}(p_{i,k}) = \frac{H(p_{i,k})}{-\log(nm)} \in [0, 1]$ of the flip distribution. We aim to find a value T such that $\bar{H}(p_{i,k}) = \bar{H}_{\text{target}}$. The relationship between the normalized entropy and the temperature is non-trivial and depends on the gradients of g . Using Equation 4, we can write the normalized entropy equation as a function of the temperature T , we solve numerically for T by minimizing by gradient descent the penalty. In practice, to select an appropriate target \bar{H}_{target} , we find it enough to test a few values in $\bar{H}_{\text{target}} \in \{0.4, 0.6, 0.8\}$. One can note that setting $\bar{H}_{\text{target}} = 1$ corresponds to completely random mutations.

4 EXPERIMENTS

4.1 SETTINGS

We conduct experiments on different benchmarks to assess the performance of our new emitter. Namely, we experiment on three challenging applications: illuminating a family of proteins, generating diverse images by working directly on the discretized pixel space and illuminating the discrete latent space of a generative model.

Design in the Ras Protein Family As proteins are involved in numerous biological phenomena, including every task of cellular life, designing them has been a longstanding issue for biologists. The simplest way to describe a protein is through its primary structure: a natural protein is a sequence of 20 basic chemical elements called amino-acids. Designing a protein is generally framed as finding the sequence of amino acids that maximizes a score representing a property of a protein, such as its ability of binding to a virus or its stability. It is a tedious task as the search space grows exponentially with the length of the protein of interest, for instance designing a protein of 100 residues operates over a space of 20^{100} elements. In our experiment, we aim to redesign members of the "Rat sarcoma virus" (*Ras*) protein family (extracted from PFAM database (Mistry et al., 2021)), involved in cellular signal transduction and plays a role in cancer development Bos (1989). Following a common procedure in computational protein design Huang et al. (2016), we aim to redesign its members to maximize their stability. While many methods use biophysical simulations to estimate this stability Desmet et al. (1992), we follow recent trends Anishchenko et al. (2021b) and use a pre-trained neural network to approximate it.

Namely, for a given sequence, our objective function is defined as the log-likelihood provided by a pre-trained protein language model: ESM2 Lin et al. (2022); Rives et al. (2021). Indeed, it was shown that this quantity is correlated to the stability of a protein and can even be used to predict the effect of these mutations Hopf et al. (2017); Riesselman et al. (2018). We use an unsupervised procedure to create descriptors similar to AURORA Grillotti & Cully (2022): we sample 1,000,000 proteins from the Uniref50¹ database, compute their 640-dimensional embeddings using ESM2 at the 20th layer and perform a PCA on those embeddings to extract 5 components. Finally we use a K -Means algorithm with $K = 30000$ on these projected embeddings to create the tessellation with 30,000 cells. The descriptors of a sequence is the corresponding projected embedding of this sequence and is mapped to its corresponding cell by finding the closest centroid. In practice, we use the 120,000 elements of this family at initialization.

¹<https://www.uniprot.org/help/uniref>

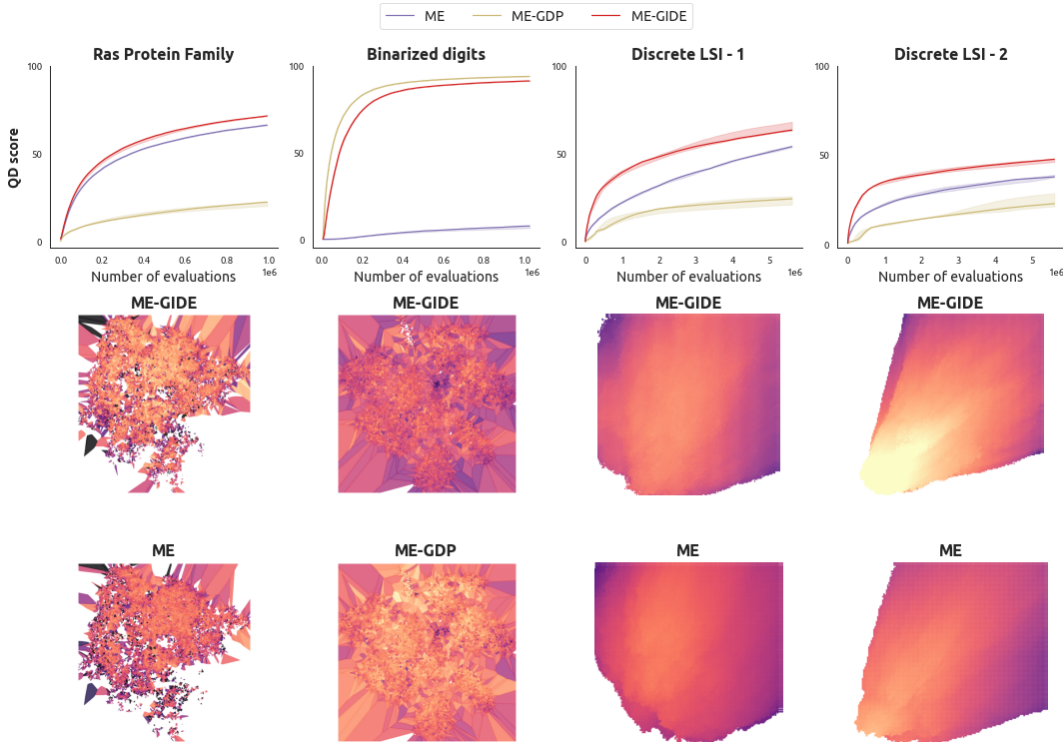


Figure 2: **(Top)** QD-score evolution on the different domains (median and interquartile range over 5 seeds). ME-GIDE outperforms MAP-ELITES on every experiment and ME-GDP on every experiment except for Binarized Digits. Only the solution corresponding to the best set of hyperparameters are shown for MAP-ELITES and ME-GDP. **(Middle)** Final repertoire of ME-GIDE for one seed. A yellow color corresponds to a high fitness, whereas darker colours correspond to lower fitnesses. **(Bottom)** Repertoire of the best performing baseline.

Binarized Digits To further illustrate the interest of our method, we design an experiment on binary data consisting in generating diverse MNIST digits with QD methods. Given our fitness and descriptor functions, we aim to find out how QD methods compare in generating diverse digits by directly searching over the image space. We define the fitness using an energy-based model trained on discrete MNIST data. Following Grathwohl et al. (2021), we train a restricted Boltzmann machine (RBM) Hinton (2012) on the binarized 28x28 images. To obtain descriptors we define a diversity space by embedding the MNIST dataset into the hidden layer of the RBM and then by computing a PCA over these embeddings. Thus, descriptors are defined as projections over the top- d components of this PCA with $d = 20$. We again finally use a K -Means algorithm with $K = 10000$ on these projected embeddings to create the tessellation with 10,000 cells. We display in Appendix D how this descriptor space spans the different digits’ classes. In our experiments, we initialize our images with a randomly uniform distribution.

Discrete Latent Space Illumination (LSI) Recent works Fontaine et al. (2021), Fontaine & Nikolaidis (2021) have proposed the latent space illumination of generative models (LSI) problem as a benchmark for QD methods. It consists in searching the latent space of a large generative model for diverse and relevant latent codes. The diversity of the sampled images being a major concern for generative models, it perfectly aligns with the QD methods design. For instance, Fontaine & Nikolaidis (2021) uses QD methods to illuminate the latent space of a StyleGAN, they leverage CLIP Radford et al. (2021) prompts to create both a fitness function and a descriptor space and to generate diverse images that satisfy some prompts. Since our work is focused on discrete variables, we instead chose to illuminate the latent space of a discrete VAE. Specifically, we choose to explore the latent space of a vector-quantized auto-encoder (VQ-VAE). Similarly to previous work, we use CLIP to evaluate the quality of the VQ-VAE-decoded latents as well as their diversity. We have experimented two sets of prompts. The first one aims to generate labrador dog images, the fitness is

defined with the prompt "A labrador" while the descriptor functions are defined as "A white puppy" and "A dog with long hair". The second set of prompt searches for truck images with fitness prompt "A truck" and descriptor prompts "A red truck", "A blue truck". The aim of this experiment is to find a large set of latent codes which - once decoded - yield diverse images with high scores according to the CLIP evaluation. We refer the reader to the Appendix C for the VQ-VAE model's details, and the training procedure. We want to stress that our goal is only to explore the latent space of the VQ-VAE in the light of the CLIP score and descriptors functions, we do not aim to compete with image generation procedure nor with text-to-image approaches.




Figure 3: *Ras* proteins designed by ME-GIDE, visualized using the prediction of AlphaFold. One can see diverse patterns in the secondary structure, while still retaining high estimated stability.

ing $1e6$ candidates on the protein and binarized digits experiment and $5e6$ evaluations for the discrete LSI experiment. We provide more details about the settings in Appendix A.

Experiment Setup In addition to our ME-GIDE algorithm, we compare with two main baselines. The first baseline, MAP-ELITES stands for the original MAP-ELITES algorithm with completely random mutations at each step. Searching via random mutations is the most popular solution when no prior knowledge is assumed over the structure of a solution. We also consider additional variation operators such as crossover. We also compare ME-GIDE to a natural baseline: gradient descent with projection (ME-GDP). Inspired by Fast SeqProp Linder & Seelig (2020), ME-GDP uses gradients following Objective and Measure Gradient MAP-ELITES via Gradient Arborecence (OMG-MEGA) Fontaine & Nikolaidis (2021) and projects the mutants back to the discrete set. We implement every method based on Jax using the QDax open source library Lim et al. (2022).

We run every experiment over 5 different seeds for each method. We use a repertoire with 30,000 centroids for the protein experiment, 10,000 centroids for the binarized digits experiments and 40,000 centroids for the discrete LSI experiments. We stop the illumination after evaluating

4.2 RESULTS AND ANALYSIS

Figure 2 summarizes the results of our experiments, we display the median and the interquartile range over 5 seeds for every benchmark. In the Ras and Discrete LSI experiments, ME-GIDE outperforms other methods, since it converges in fewer iteration to higher performances. For clarity, we only plot the results corresponding for the best set of hyperparameters for the baselines MAP-ELITES and ME-GDP. We use a Wilcoxon signed-rank test Wilcoxon (1992) to compare the distribution of the scores obtained over different seeds where the null hypothesis is that the obtained scores have the same median. It shows that ME-GIDE constantly outperforms MAP-ELITES in the QD score with p-values lower than 0.05 in every experiment and outperforms ME-GDP on every experiment except Binarized Digits.

Our target entropy procedure allows an easier choice of hyperparameters for ME-GIDE. Indeed, in our experiment, we only try for three values: 0.4, 0.6 and 0.8. While we display the best results for clarity on Figure 2, the conclusion are similar for values in the range $[0.4, 0.8]$ as detailed in Appendix B.

Interestingly, ME-GDP under-performs compared to purely random mutations for any set of hyperparameters on both protein and VQ-VAE experiments. This means performing gradient descent

similarly as in the continuous case is dangerous when using discrete variables. We validate that the estimation of $g(x^{(ik)}) - g(x)$ using \tilde{d}_{ik} is sensible by computing the correlation for some elements of each domain between those two quantities. For instance, for the discrete LSI experiment, we find an average correlation $\rho = 0.55$, which means that even though the step between $x^{(ik)}$ and x might be large, gradient information is indeed relevant. We give more details about this analysis in Appendix E.

We display some samples of the final repertoire obtained by ME-GIDE on Figure 3. To visualize the 3D structure of those proteins, we use AlphaFold2 Jumper et al. (2021b). One can see that ME-GIDE manages to extract proteins with various secondary and tertiary structures and that those proteins have high confidence score according to AlphaFold. We also conduct further validation on the protein experiment. We subsample our repertoire by recomputing larger cells to obtain a repertoire with 300 cells. Then, we insert the original repertoire into this smaller one to obtain the most fit solutions in each region. We firstly analyze the diversity in the sequence space, defined as the average edit distance between obtained solutions. We find that ME-GIDE reaches an average distance of 106.9 whereas MAP-ELITES gets 97.2. We also use *SAPred* Moffat & Jones (2021) to predict the secondary structure of each protein and use the average edit distance between secondary structures. ME-GIDE outperforms MAP-ELITES, obtaining an average distance of 44.2 against 34.9, and finds more diverse structures while also reaching an overall higher average fitness.

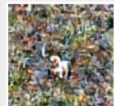

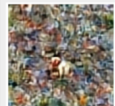

ME-GIDE samples vs ImageNet images			
Image	CLIP-score	Image	CLIP-score
	70.2		66.2
	70.9		66.7

Figure 5: Samples found by ME-GIDE on the "Labrador" prompt against highest scoring images from ImageNet dataset. While the images found by ME-GIDE are not natural images, they achieve a higher score to be a "labrador" according to CLIP.

latent codes in the latent space of a discrete VAE, which is in general bypassed by an additional auto-regressive model that learns to sample meaningful latent codes.

5 RELATED WORKS

Discrete Optimization Discrete optimization Parker & Rardin (2014) is a longstanding branch of optimization where some of the variables or constraints are discrete. It is also named combinatorial optimization as the set of feasible solutions is generally finite but evaluation every solution is irrealizable. A first class of algorithms search for an exact solution: in particular, Integer Linear Proramming Schrijver (1998) has been the main subject of interest with exact methods such as cutting plan Gilmore & Gomory (1961) or branch and bound Land & Doig (2010). On the other hand, evolutionary algorithms have been very popular as they allow for more flexibility and can be stopped at any iteration to get the best solution explored so far Dorigo et al. (1999). Most of those




Figure 4: Samples from ME-GIDE final repertoire on the Binarized Digits experiment. We down-sample the 40,000 elements in the repertoire by clustering its centroids to 16 and display the element in each cell.

On Figure 4, we first display some samples found by ME-GIDE on the binarized MNIST experiment. One can see that ME-GIDE is able to find diverse images that resemble MNIST images and covers every digit. We also visualize on Figure 5 top-performing images from the repertoire learned by ME-GIDE as well as natural images from the ImageNet dataset. Images found by ME-GIDE cannot be considered to match the prompt from a human point of view although a small shape of a dog appears in the center of the image. It should still be noted that these images have higher CLIP scores than natural images that humans would score as perfectly aligned with the score prompt. In a sense, ME-GIDE manages to "fool" CLIP. Our work emulates the findings of Nguyen et al. (2015a) where Evolutionary Algorithms manage to "fool" a classifier. Our study highlights the difficulty of directly sampling relevant latent codes.

algorithms either incorporate this knowledge by mutating only to feasible solutions or use a relaxation and then project back to the feasible set Lin & Hajela (1992). For instance, for protein design, exact methods have been proposed for specific biophysical cost functions related to the potential energy Desmet et al. (1992) of the protein or its binding affinity Ojewole et al. (2018). However, due to the limitations in flexibility of the previous methods, the most popular design tool, incorporated in the popular *Rosetta* library, is based on simulated annealing Huang et al. (2011). Several work tried to leverage recent breakthroughs of deep learning methods tailored to proteins in the protein design, such as *AlphaDesign* Jendrusch et al. (2021) or *TrDesign* Norn et al. (2021). To the best of our knowledge, our work is the first to propose QD as a solution for protein design.

Quality Diversity Optimization Searching for novelty instead of quality Lehman & Stanley (2011a) was the first work formulating diversity as an end in itself. It was further refined by introducing a notion of local quality with local competition Lehman & Stanley (2011b). MAP-ELITES Mouret & Clune (2015) refined the notion of diversity by introducing the notion of repertoire over a set of descriptors. Further improvement were made on the design the descriptor space such as using a Voronoi tessellation Vassiliades et al. (2017) or use unsupervised descriptors Cully (2019); Grillotti & Cully (2022) as defining tailored descriptors can be tedious for some tasks. More efficient ways to illuminate this space have been proposed, such as more efficient mutation operators Vassiliades & Mouret (2018) or covariance matrix adaptation Fontaine et al. (2020). Closest to this work is Differentiable QD Fontaine & Nikolaidis (2021) where gradients are used to directly update over the continuous variables.

Application specific methods were developed to apply QD algorithms to noisy domains Flageat & Cully (2020) or when multiple objectives are at stake Pierrot et al. (2022b). Reinforcement Learning is one of the most popular application of QD Arulkumaran et al. (2019) and some methods try to incorporate diversity directly in RL algorithms Nilsson & Cully (2021); Pierrot et al. (2022a). QD methods have already been applied to Latent Space Exploration, one of the domains we experiment on. The earliest work was named Innovation Engines Nguyen et al. (2015b) where authors try to generate diverse images using a QD approach over ImageNet. Later Latent Space Illumination has been introduced to generate game levels Schrum et al. (2020) or images Fontaine & Nikolaidis (2021). Our work is also applicable to LSI in the case of a discrete latent space and is the first work of its kind to the best of authors’ knowledge.

6 CONCLUSION

We introduced ME-GIDE, a genetic algorithm that leverages gradient information over discrete variables for Quality Diversity optimization and show it outperforms classical baselines. Our entropy-based guidance allows to ease the search for good hyperparameters which can be tedious for some applications. Our experiment on protein design is the first using QD methods to the best of author’s knowledge and this opens a way for more applications in the future. Indeed, we believe that generating not only fit but also diverse proteins can be useful as the fitness given by the model does not always translate to good properties in the lab.

Our goal was not to propose a novel image generation technique but while the results of our LSI experiment are interpretable and ME-GIDE finds images with higher fitnesses than realistic images, further research can be done to better understand and exploit these results. Our method can be used to create adversarial examples for generative models using discrete latent spaces, that are becoming more and more popular. On the other hand, to generate more realistic images, we could constrain the search to a few latent codes and use a recurrent architecture to generate a more realistic prior, at the cost of potentially facing vanishing or exploding gradients. We leave these open questions for future works.

As our method involves computing gradients of the fitness and descriptor functions, we expect it to be intractable for functions involving large neural networks, such as AlphaFold. Using a surrogate model to guide the search while keeping a large model for evaluation or validation to handle those cases could be a good research direction.

REFERENCES

- Ivan Anishchenko, Samuel J. Pellock, Tamuka M. Chidyausiku, Theresa A. Ramelot, Sergey Ovchinnikov, Jingzhou Hao, Khushboo Bafna, Christoffer Norn, Alex Kang, Asim K. Bera, Frank DiMaio, Lauren Carter, Cameron M. Chow, Gaetano T. Montelione, and David Baker. De novo protein design by deep network hallucination. *Nature*, 600(7889):547–552, Dec 2021a. ISSN 1476-4687. doi: 10.1038/s41586-021-04184-w. URL <https://doi.org/10.1038/s41586-021-04184-w>.
- Ivan Anishchenko, Samuel J Pellock, Tamuka M Chidyausiku, Theresa A Ramelot, Sergey Ovchinnikov, Jingzhou Hao, Khushboo Bafna, Christoffer Norn, Alex Kang, Asim K Bera, et al. De novo protein design by deep network hallucination. *Nature*, 600(7889):547–552, 2021b.
- Kai Arulkumaran, Antoine Cully, and Julian Togelius. Alphastar: An evolutionary computation perspective. In *Proceedings of the genetic and evolutionary computation conference companion*, pp. 314–315, 2019.
- Johannes L Bos. Ras oncogenes in human cancer: a review. *Cancer research*, 49(17):4682–4689, 1989.
- Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret. Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*, pp. 109–135. Springer, 2021.
- Antoine Cully. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 81–89, 2019.
- Antoine Cully and Yiannis Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259, 2017.
- Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- Johan Desmet, Marc De Maeyer, Bart Hazes, and Ignace Lasters. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature*, 356(6369):539–542, 1992.
- Marco Dorigo, Gianni Di Caro, and Luca M Gambardella. Ant algorithms for discrete optimization. *Artificial life*, 5(2):137–172, 1999.
- Manon Flageat and Antoine Cully. Fast and stable map-elites in noisy domains using deep grids. *arXiv preprint arXiv:2006.14253*, 2020.
- Matthew Fontaine and Stefanos Nikolaidis. Differentiable quality diversity. *Advances in Neural Information Processing Systems*, 34:10040–10052, 2021.
- Matthew C Fontaine, Julian Togelius, Stefanos Nikolaidis, and Amy K Hoover. Covariance matrix adaptation for the rapid illumination of behavior space. In *Proceedings of the 2020 genetic and evolutionary computation conference*, pp. 94–102, 2020.
- Matthew C Fontaine, Ruilin Liu, Ahmed Khalifa, Jignesh Modi, Julian Togelius, Amy K Hoover, and Stefanos Nikolaidis. Illuminating mario scenes in the latent space of a generative adversarial network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 5922–5930, 2021.
- Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. *arXiv preprint arXiv:2203.13131*, 2022.
- Paul C Gilmore and Ralph E Gomory. A linear programming approach to the cutting-stock problem. *Operations research*, 9(6):849–859, 1961.
- Will Grathwohl, Kevin Swersky, Milad Hashemi, David Duvenaud, and Chris Maddison. Oops i took a gradient: Scalable sampling for discrete distributions. In *International Conference on Machine Learning*, pp. 3831–3841. PMLR, 2021.

- Luca Grillotti and Antoine Cully. Unsupervised behaviour discovery with quality-diversity optimisation. *IEEE Transactions on Evolutionary Computation*, 2022.
- Geoffrey E Hinton. A practical guide to training restricted boltzmann machines. In *Neural networks: Tricks of the trade*, pp. 599–619. Springer, 2012.
- Thomas A Hopf, John B Ingraham, Frank J Poelwijk, Charlotta PI Schärfe, Michael Springer, Chris Sander, and Debora S Marks. Mutation effects predicted from sequence co-variation. *Nature biotechnology*, 35(2):128–135, 2017.
- Po-Ssu Huang, Yih-En Andrew Ban, Florian Richter, Ingemar Andre, Robert Vernon, William R Schief, and David Baker. Rosettaremodel: a generalized framework for flexible backbone protein design. *PloS one*, 6(8):e24109, 2011.
- Po-Ssu Huang, Scott E Boyken, and David Baker. The coming of age of de novo protein design. *Nature*, 537(7620):320–327, 2016.
- Michael Jendrusch, Jan O Korbelt, and S Kashif Sadiq. Alphadesign: A de novo protein design framework based on alphafold. *bioRxiv*, 2021.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, Aug 2021a. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2. URL <https://doi.org/10.1038/s41586-021-03819-2>.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021b.
- Ailsa H Land and Alison G Doig. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pp. 105–132. Springer, 2010.
- Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011a.
- Joel Lehman and Kenneth O Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 211–218, 2011b.
- Bryan Lim, Maxime Allard, Luca Grillotti, and Antoine Cully. Accelerated quality-diversity for robotics through massive parallelism. *arXiv preprint arXiv:2202.01258*, 2022.
- C-Y Lin and Prabhat Hajela. Genetic algorithms in optimization problems with discrete and integer design variables. *Engineering optimization*, 19(4):309–327, 1992.
- Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Sal Candido, and Alexander Rives. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022. doi: 10.1101/2022.07.20.500902. URL <https://www.biorxiv.org/content/early/2022/07/21/2022.07.20.500902>.
- Johannes Linder and Georg Seelig. Fast differentiable dna and protein sequence optimization for molecular design. *arXiv preprint arXiv:2005.11275*, 2020.
- Jaina Mistry, Sara Chuguransky, Lowri Williams, Matloob Qureshi, Gustavo A Salazar, Erik LL Sonnhammer, Silvio CE Tosatto, Lisanna Paladin, Shriya Raj, Lorna J Richardson, et al. Pfam: The protein families database in 2021. *Nucleic acids research*, 49(D1):D412–D419, 2021.

- Lewis Moffat and David T Jones. Increasing the accuracy of single sequence prediction methods using a deep semi-supervised learning framework. *Bioinformatics*, 37(21):3744–3751, 2021.
- Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. 2015.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 427–436, 2015a.
- Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pp. 959–966, 2015b.
- Olle Nilsson and Antoine Cully. Policy gradient assisted map-elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 866–875, 2021.
- Christoffer Norn, Basile IM Wicky, David Juergens, Sirui Liu, David Kim, Doug Tischer, Brian Koepnick, Ivan Anishchenko, Foldit Players, David Baker, et al. Protein sequence design by conformational landscape optimization. *Proceedings of the National Academy of Sciences*, 118(11):e2017228118, 2021.
- Adegoke A Ojewole, Jonathan D Jou, Vance G Fowler, and Bruce R Donald. Bbk*(branch and bound over k*): A provable and efficient ensemble-based protein design algorithm to optimize stability and binding affinity over large sequence spaces. *Journal of Computational Biology*, 25(7):726–739, 2018.
- R Gary Parker and Ronald L Rardin. *Discrete optimization*. Elsevier, 2014.
- Thomas Pierrot, Valentin Macé, Felix Chalumeau, Arthur Flajolet, Geoffrey Cideron, Karim Beguir, Antoine Cully, Olivier Sigaud, and Nicolas Perrin-Gilbert. Diversity policy gradient for sample efficient quality-diversity optimization. In *ICLR Workshop on Agent Learning in Open-Endedness*, 2022a.
- Thomas Pierrot, Guillaume Richard, Karim Beguir, and Antoine Cully. Multi-objective quality diversity optimization. *arXiv preprint arXiv:2202.03057*, 2022b.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.
- Adam J Riesselman, John B Ingraham, and Debora S Marks. Deep generative models of genetic variation capture the effects of mutations. *Nature methods*, 15(10):816–822, 2018.
- Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021.
- Jason Tyler Rolfe. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- Alexander Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- Jacob Schrum, Vanessa Volz, and Sebastian Risi. Cppn2gan: Combining compositional pattern producing networks and gans for large-scale pattern generation. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pp. 139–147, 2020.

Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

Vassiliis Vassiliades and Jean-Baptiste Mouret. Discovering the elite hypervolume by leveraging interspecies correlation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 149–156, 2018.

Vassilis Vassiliades, Konstantinos Chatzilygeroudis, and Jean-Baptiste Mouret. Using centroidal voronoi tessellations to scale up the multidimensional archive of phenotypic elites algorithm. *IEEE Transactions on Evolutionary Computation*, 22(4):623–630, 2017.

Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pp. 196–202. Springer, 1992.