Victor Costa University of Coimbra, CISUC, DEI Coimbra, Portugal vfc@dei.uc.pt

João Correia University of Coimbra, CISUC, DEI Coimbra, Portugal jncor@dei.uc.pt

## ABSTRACT

Generative adversarial networks (GANs) achieved relevant advances in the field of generative algorithms, presenting high-quality results mainly in the context of images. However, GANs are hard to train, and several aspects of the model should be previously designed by hand to ensure training success. In this context, evolutionary algorithms such as COEGAN were proposed to solve the challenges in GAN training. Nevertheless, the lack of diversity and premature optimization can be found in some of these solutions. We propose in this paper the application of a quality-diversity algorithm in the evolution of GANs. The solution is based on the Novelty Search with Local Competition (NSLC) algorithm, adapting the concepts used in COEGAN to this new proposal. We compare our proposal with the original CO-EGAN model and with an alternative version using a global competition approach. The experimental results evidenced that our proposal increases the diversity of the discovered solutions and leverage the performance of the models found by the algorithm. Furthermore, the global competition approach was able to consistently find better models for GANs.

#### CCS CONCEPTS

 $\bullet$  Computing methodologies  $\rightarrow$  Genetic algorithms; Neural networks;

### **KEYWORDS**

neuroevolution, coevolution, generative adversarial networks, quality diversity

#### ACM Reference Format:

Victor Costa, Nuno Lourenço, João Correia, and Penousal Machado. 2020. Exploring the Evolution of GANs through Quality Diversity. In *Genetic and Evolutionary Computation Conference (GECCO* 

GECCO '20, July 8–12, 2020, Cancún, Mexico © 2020 Association for Computing Machinery. ACM ISBN 978-1-4503-7128-5/20/07...\$15.00

https://doi.org/10.1145/3377930.3389824

Nuno Lourenço University of Coimbra, CISUC, DEI Coimbra, Portugal naml@dei.uc.pt

Penousal Machado University of Coimbra, CISUC, DEI Coimbra, Portugal machado@dei.uc.pt

'20), July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3377930.3389824

#### **1 INTRODUCTION**

Generative Adversarial Networks (GANs) [9] are an adversarial model that gained a lot of relevance in recent years, mainly by the success in generative tasks. Even though GANs can be applied in several contexts such as image, video, sound, and text, the GAN model is popular for the impressive results concerning the quality of created samples in the context of images. The GAN model consists of two neural networks: one generator and one discriminator, trained in an adversarial way. A successful GAN training produces strong generative and discriminative components.

Despite the compelling results, the training of GANs is challenging and is frequently affected by the presence of problems such as the vanishing gradient and the mode collapse [3, 7]. Improvements over the original model were introduced to handle these issues, but they are still a problem [2, 10, 24]. Alternatives loss functions such as in WGAN [2], LSGAN [17], and SN-GAN [19] were proposed to improve the model. Besides, architectural guides and strategies were developed to minimize these issues [12, 22].

Another strategy to improve the training of GANs is the application of evolutionary algorithms. In this context, recent solutions were designed to improve the training process and the quality of the outcome [1, 4, 5, 8, 27, 28]. These proposals incorporate mechanisms such as neuroevolution, coevolution, and spatial coevolution, making use of evolutionary pressure to achieve efficient models.

Coevolutionary Generative Adversarial Networks (CO-EGAN) [4, 5] is a solution inspired on NEAT [25] and Deep-NEAT [18] that applies a coevolution model to evolve GANs. Experimental results show that the model provides a more reliable GAN training when compared to regular GANs in equivalent scenarios. However, the lack of diversity evidenced in the experimental evaluation affects the quality of the results, leaving space for improvement of the algorithm that drives the evolutionary process. Thus, we study in this work a mechanism of novelty to be applied in COEGAN in order to improve the exploration of solutions. Quality Diversity (QD) algorithms are a class of solutions that can be used to enhance the population and produce a diversity of efficient individuals [21]. At the best of our knowledge, no proposed

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

solutions combining evolutionary algorithms and GANs use novelty search or QD mechanisms in their approaches.

In this paper, we propose a new model combining concepts used on COEGAN and a quality diversity algorithm for guiding the evolution of GANs. Therefore, instead of strategies such as the speciation based on NEAT to support evolution, we propose the use of Novelty Search with Local Competition (NSLC) [14], a quality diversity algorithm that uses mechanisms of novelty in the search for efficient solutions. We aim to improve the exploration of the search space and achieve better models for generators and discriminators.

To validate our proposal, experiments were conducted using the MNIST [13] and CelebA [15] datasets. We compare the results between the original COEGAN approach and two variations of our proposal: COEGAN with NSLC and with an alternative using a global competition strategy. The results evidenced that the exploration of solutions was improved with the QD algorithm, leading to the discovery of more efficient models for GANs. Besides, the global competition version achieved the best results concerning the quality of samples created by generators.

The remainder of this paper is organized as follows: Section 2 introduces concepts related to GANs, also describing the works associated with evolutionary algorithms; Section 3 presents our approach to use a quality diversity algorithm in the evolution of GANs; Section 4 displays the experimental results of this approach; finally, Section 5 presents our conclusions.

### 2 BACKGROUND AND RELATED WORKS

Generative Adversarial Networks (GANs) represent an adversarial model composed of two neural networks: a generator and a discriminator. The discriminator receives a dataset as input and has to distinguish between samples of this dataset and fake samples. The generator is responsible for producing synthetic data in order to fool the discriminator. As the training progresses, both the generator and the discriminator improve their tasks, resulting in strong generative and discriminative components at the end of a successful training.

The regular GAN training uses backpropagation and gradient descent in both neural networks. Thus, the loss function of the discriminator is defined as follows:

$$J^{(D)}(D,G) = -\mathbb{E}_{x \sim p_{data}}[\log D(x)] - \mathbb{E}_{z \sim p_z}[\log(1 - D(G(z)))].$$
(1)

For the generator, the non-saturating version of the loss function is defined by:

$$J^{(G)}(G) = -\mathbb{E}_{z \sim p_z}[\log(D(G(z)))].$$
 (2)

In Eq. (1),  $p_{data}$  represents the dataset used as input to the discriminator. In Eq. (1) and Eq. (2), z is the latent space used as input to the generator,  $p_z$  is the latent distribution, G is the generator, and D represents the discriminator.

GANs are hard to train and a trial-and-error approach is frequently used to get consistent results. The equilibrium of forces between the discriminator and the generator is frequently the cause of problems in training. In the vanishing gradient problem, the discriminator or generator is so powerful that it becomes almost perfect in its task, leading to the stagnation of training progress. The mode collapse occurs when the generator captures only a small fraction of the input distribution, limiting the diversity of produced samples.

In GANs, the Fréchet Inception Distance (FID) [11] is often used to evaluate the performance of the generators. The FID metric uses the Inception Net [26] (trained on ImageNet [23]) to transform images into a feature space, which is interpreted as a continuous multivariate Gaussian. This process is applied to samples from the input dataset and synthetic samples created by the generator. The mean and covariance of the two resulting Gaussians are estimated and the Fréchet distance between them is given by:

$$FID(x,g) = ||\mu_x - \mu_g||_2^2 + Tr(\Sigma_x + \Sigma_g - 2(\Sigma_x \Sigma_g)^{1/2}), (3)$$

with  $\mu_x$ ,  $\Sigma_x$ ,  $\mu_g$ , and  $\Sigma_g$  representing the mean and covariance estimated for the input dataset x and fake samples g, respectively. This metric is capable of quantifying the quality and diversity of the generative model.

The use of evolutionary algorithms to train and evolve GANs was recently proposed [1, 4, 5, 8, 27, 28]. The solutions present a diverse set of strategies to not only overcome common GAN problems but also to provide better quality on the produced samples.

E-GAN [28] evolve GANs using a variation operator that switches the loss function of the generator through generations. In this case, a single-fixed discriminator is used as the adversarial for the population of generators, with the former using a fixed architecture and loss function, and the latter varying only the loss function. The architectures of the generator and the discriminator are based on DCGAN [22]. In [8], Pareto set approximation was used in a neuroevolution algorithm to evolve GANs. In this case, the architectures of the networks are dynamic and change according to the variation operators. Lipizzaner [1] uses spatial coevolution to train GANs. However, the networks of the discriminator and generator are fixed and only the internal parameters (e.g., weights) change through evolution. A further improvement over Lipizzaner, called Mustangs [27], applies the E-GAN dynamic loss function to the algorithm while keeping the same spatial coevolution strategy of Lipizzaner.

COEGAN [4, 5] was inspired by NEAT [25] and Deep-NEAT [18] to design an evolutionary algorithm for GANs, using mechanisms such as speciation to protect innovation during the evolution. In COEGAN, the fitness used for generators is based on the FID score (Eq. (3)). The use of FID was designed to put selection pressure in generators and guide the evolution of the population in producing better samples. For discriminators, fitness is based on the loss function represented by Eq. (1).

Nondominated Sorting Genetic Algorithm II (NSGA-II) [6] is a well-known solution in the class of Multi-Objective Evolutionary Algorithms (MOEAs) that uses an elitist method to implement a Pareto-based search approach. In NSGA-II,

an algorithm to sort solutions and determine nondominated fronts is used on the generation of the next populations. Besides, a crowding-distance computation is used as a second criterion to prioritize solutions in less explored spaces. This algorithm also uses an archive to keep the previously explored solutions and improve diversity. Individuals are usually inserted into the archive using a probabilistic approach.

Quality Diversity (QD) algorithms are a family of evolutionary algorithms aiming to find a diversity of viable solutions for the target problem [21]. Novelty Search with Local Competition (NSLC) [14] uses a Pareto-based MOEA, such as NSGA-II, to promote the quality and diversity of solutions. NSLC uses as objectives the quality and novelty of individuals according to a local neighborhood. Therefore, when combined with NSGA-II, NSLC does not use the crowding-distance mechanism because the novelty criterion already produces the desired diversity.

MAP-Elites [20] is another solution in the class of QD algorithms. In MAP-Elites, an archive of phenotypes is kept in order to explore the diversity of high-performing solutions. Thus, at each step, an item of the archive is chosen to produce the offspring. Then, the performance is calculated and the newly generated individual is placed into the archive in the position determined by the feature space, replacing older individuals in case of better performance. The MAP-Elites algorithm was also considered to be the subject of this research but the cost to explore and maintain an archive of high-dimensional neural networks was considered too high.

#### 3 METHODS

We describe in this section our proposal to use a Quality Diversity (QD) algorithm to evolve GANs. This new model is based on the original COEGAN proposal [4, 5], adapted to be guided by a different evolutionary algorithm. Thus, first we introduce the fundamentals of the COEGAN model. Then, we describe our approach to applying the QD algorithm with COEGAN.

#### 3.1 COEGAN

In COEGAN, the genome consists of an array of genes that are directly mapped into sequential layers of a neural network, forming the phenotype of individuals. This approach was inspired by NEAT [25] and its extension DeepNEAT [18]. The genes describe either a linear, convolution or transpose convolution layer (also known as deconvolution layer), depending on the type of individuals. Generators allow linear and transpose convolution layers, and discriminators allow linear and convolution layers. In addition, each gene holds internal parameters specific to the type of layer. Parameters such as the activation function, the number of output features, and the number of output channels are subject to variation operators.

The variation operators are based on mutation, with the possibilities to add, remove or change genes. The addition operator inserts a new gene into the genome. This new layer is randomly drawn from a set of possible layers: linear and convolution for discriminators; linear and transpose convolution for generators. The removal operation randomly chooses an existing gene and excludes it from the genome. The change operation modifies the internal attributes such as the activation function of an existing layer. In this case, the activation function is randomly chosen from the set: ReLU, LeakyReLU, ELU, Sigmoid, and Tanh. For dense and convolution layers, the number of output features and channels can also be mutated, respectively. The mutation of these attributes follows a uniform distribution, delimited by a predefined range.

In COEGAN, the weights of parents are transferred to the children whenever it is possible [4, 5]. This mechanism of transference ensures that the information achieved on training in previous generations is kept during the whole evolution process. Therefore, not only the final models achieved by COEGAN are important but also the entire process used in the discovery of them.

Competitive coevolution was used to model the COEGAN algorithm, creating two independent subpopulations of generators and discriminators. Thus, the evaluation phase matches generators and discriminators for training and evaluation. The *all vs. all* approach is used to pair each generator with each discriminator in the regular GAN training algorithm.

A fitness sharing strategy is used to protect the species and promote innovation. Therefore, each subpopulation of generators and discriminators is divided into species. The speciation mechanism was inspired by NEAT and aims to promote innovation in the population, ensuring that recently modified individuals will be trained for enough generations to be comparable to older individuals with respect to the fitness value. A similarity function comparing pairs of genomes is used to group individuals into species.

At the evaluation phase, the fitness used for discriminators is based on the loss function of the discriminator in the original GAN model, represented by Eq. (1). For generators, the fitness is the FID score, represented by Eq. (3).

#### 3.2 Quality Diversity in COEGAN

In this work, we propose a new evolutionary algorithm to guide the COEGAN training. Therefore, we replace the NEAT-based evolutionary algorithm used in COEGAN with an approach based on Novelty Search with Local Competition (NSLC) [14]. As originally proposed in [14], we use Nondominated Sorting Genetic Algorithm II (NSGA-II) [6] as the Multi-Objective Evolutionary Algorithm (MOEA) for NSLC.

We kept some aspects used in COEGAN regarding the genotype representation. Thus, in our proposal, the genome and variation operators are the same originally used in CO-EGAN, described in Section 3.1. The competitive coevolution model is also the same used in COEGAN. Following we describe the differences and new aspects of the model proposed in this work.

We designed the pairing between individuals at the evaluation phase according to the *all vs. all* coevolution model used in COEGAN. However, as NSGA-II uses an elitist approach to select individuals to form the next generation, both the current population and the derived offspring should be evaluated. For this, we match generators from the current population with discriminators from the offspring, and discriminators from the current population. This ensures the progress of all individuals through generations, making it possible to properly select them when targeting for quality and diversity. The main drawback here is that now the double amount of individuals should be evaluated, increasing the execution cost of the algorithm. However, it is important to note that we keep the same number of training steps for each individual when compared to the original COEGAN approach.

Tournament is applied to select individuals for reproduction. As proposed by the NSGA-II algorithm, a dominance operator is used to determine the result of the tournament between a set of individuals. In this research, we use the constrained version of the operator, ensuring that the population does not deviate too much from the objective defined by the fitness function. Thus, we use not only the concept of dominance but also the feasibility of solutions.

NSGA-II designs the dominance operator using the ranking of solutions determined by the nondominated sorting algorithm, aiming to obtain solutions in the Pareto-optimal front. The feasibility concept ensures that, when comparing two solutions  $s_i$  and  $s_j$ , the fitness function meets the constraint  $f(s_j) < 2f(s_i)$ , otherwise  $s_j$  is considered unfeasible. In summary, the solution  $s_i$  constrained-dominates a solution  $s_j$  when  $s_i$  is feasible and  $s_j$  is not, or both solutions are feasible and  $s_i$  dominates  $s_j$ . Note that one of the solutions will always be feasible, i.e., the case that both solutions are unfeasible is not possible.

The definition of the neighborhood is paramount to the NSLC algorithm, being used in both the quality and diversity criteria. In order to determine the neighbors of each individual, we use the distance between the architectures of the neural networks of individuals, which is directly defined by the similarity between genomes. This distance is the same used originally in COEGAN to group individuals into species. Two individuals are considered equal if they have the same genome, i.e., the same sequence of genes, disregarding other characteristics like age or the number of samples currently used in the training. It is important to note that the neighborhood calculation considers not only the current population but also the archive of previous solutions. This archive is filled following a probabilistic approach, i.e., at each generation individuals are inserted into the archive with a predefined probability.

In NSLC, n nearest neighbors of an individual are selected to calculate the innovation and the competition objectives. Innovation is defined by the average distance between the individual and the neighborhood. The competition objective is defined by the number of neighbors the individual outperforms with respect to the fitness value. In our proposal, the fitness is the same used in COEGAN: Eq. (1) for discriminators and Eq. (2) for generators. The innovation criteria make it possible to better explore the available architectures characterized by the genotype representation. When combined with the strategy used to calculate the competition score, different niches can be efficiently explored to leverage the search space. Compared to the original COEGAN approach, we expect to improve the diversity of solutions and eventually find better results concerning the FID score. In COEGAN, the number of species is fixed and previously defined, being a limitation over the capacity of innovation for individuals that need to survive through generations to show consistent performances. In CO-EGAN guided by NSLC, the exploration of the search space is improved by the novelty criterion, using the quality definition to guide the population through the objective of obtaining better solutions.

#### 4 EXPERIMENTS

To validate our proposal, we present an experimental analysis of the application of Quality Diversity in the evolution of GANs<sup>1</sup>. Therefore, we conduct experiments using MNIST in order to evidence the performance of the algorithm proposed in this work compared with the original COEGAN model. We also design experiments with an alternative version of the solution which uses a global competition mechanism, i.e., the neighborhood is not limited by a constant n and uses all individuals available. We call this version of the algorithm Novelty Search with Global Competition (NSGC), inspired by the global competition approach experimented in [14]. Therefore, we refer to the original COEGAN proposal, COEGAN trained with the NSLC algorithm and trained with NSGC as COEGAN, COEGAN+NSLC, and COEGAN+NSGC, respectively.

The FID score was used to measure the quality of the produced samples. Besides, the strategy proposed in [29] was applied to present the visual distribution of image samples, using t-SNE [16] to embed samples into a two-dimensional space. Further experiments with the CelebA dataset were made to compare our method with a non-evolutionary GAN approach in a more complex dataset.

#### 4.1 Experimental Setup

Table 1 lists the parameters used in our experiments. These parameters were selected based on preliminary experiments and the experiments presented in [4, 5]. The number of generations used in all experiments is 50. Each population of generators and discriminators contains 10 individuals. We use a probability of 30%, 10% and 10% for the add, remove and change mutations, respectively. The genome of generators and discriminators was limited to four genes, representing a network of four layers in the maximum allowed setup. This setup is sufficient to discover efficient solutions for the experiments with the MNIST dataset.

For the original COEGAN, we use 3 species in each population of generators and discriminators. For COEGAN with NSLC, the number of neighbors n is limited to 3 and the

<sup>&</sup>lt;sup>1</sup>Code available at https://github.com/vfcosta/qd-coegan.

Table	: 1:	Experimental	parameters.
-------	------	--------------	-------------

<b>Evolutionary Parameters</b>	Value
Number of generations	50
Population size (generators)	10
Population size (discriminators)	10
Add Layer rate	30%
Remove Layer rate	10%
Change Layer rate	10%
Output channels range	[32, 256]
Tournament $k_t$	2
FID samples	1024
Genome Limit	4
Species	3
Neighborhood size $n$	3
Archive probability	10%
GAN Parameters	Value
Batch size	64
Batches per generation	50
Optimizer	Adam
Learning rate	0.001

probability to insert individuals into the archive is 10%. The global version of the QD algorithm does not limit the neighborhood, using all individuals when calculating the novelty and competition values.

Figures in this section display plots with curves representing the average of the results from 15 independent executions, with a confidence interval of 95%.

#### 4.2 Results

We present in this section the results of the experimental analysis, comparing the solutions using the QD algorithm with the previously proposed COEGAN model. First, we present the results using the MNIST dataset. Then, we provide a further analysis with CelebA, a more complex dataset, comparing our proposal with a regular GAN approach.



Figure 1: Best FID Score on the MNIST dataset.

Figure 1 shows the FID score of the best individuals for each generation when training with the MNIST dataset. We

# Table 2: Average FID score of best generators after training with the MNIST dataset.

Algorithm	FID Score
COEGAN	$36.8 \pm 18.6$
COEGAN+NSLC	$35.2 \pm 12.5$
COEGAN+NSGC	$24.3 \pm 3.3$

can see that the COEGAN+NSLC solution outperforms the original COEGAN model by a small margin until half generations but has equivalent performance at the end. This effect is mostly due to the increased exploration capability given by the QD algorithm, which produces a more diverse population but causes less focused evolution of more fitted individuals. Besides, COEGAN+NSGC, the global competition variation, has better performance than COEGAN+NSLC and the original COEGAN approach. The results obtained by comparing the global and local competition versions of the algorithm are similar to results presented in [6], where the global version also achieved better fitness than NSLC.



Figure 2: Boxplot of the FID score on MNIST dataset showing the performance of best generators computed for each independent run.

We can see in Figure 2 and Table 2 that COEGAN+NSGC consistently achieved better results than the other solutions. COEGAN provides more unstable results when compared to the global approach, presenting a higher standard deviation in FID values when trained with the experimental setup described in Table 1. This effect is also present in the experiments with COEGAN+NSLC, indicating that the high diversity produced by our experimental setup affects the results with respect to the FID score. Therefore, we conclude that the diversity provided by COEGAN+NSGC is sufficient to achieve better results in our approach to train GANs.

To support our analysis, we statistically test the significance of our findings. We assume that the results do not follow a normal distribution, as the normality test (Shapiro-Wilk,  $\alpha = 0.05$ ) rejected this hypothesis for COEGAN and CO-EGAN+NSLC (p < 0.001). Then, we used a non-parametric test (Mann-Whitney U,  $\alpha = 0.05$ ) to perform a pairwise comparison between the solutions evaluated in this work. We found that the improvement of COEGAN+NSGC over COEGAN is statistically significant (p = 0.008). Moreover, the performance improvement of COEGAN+NSGC over COEGAN+NSLC is also statistically significant (p = 0.0001). We found no statistical difference between COEGAN and COEGAN+NSLC (p = 0.17). Further experiments should be performed to assess the influence of the experimental setup in these results.



Figure 3: Average FID Score on the MNIST dataset.

We also show in Figure 3 the FID scores of all individuals in the population of generators when training with MNIST. This result evidences that COEGAN+NSLC has a better exploration of the search space, increasing the diversity and leading to the discovery of not only good individuals but also less efficient solutions. The local competition approach used in NSLC has a stronger effect on the protection of innovation, as the competition calculation uses the fitness values from similar individuals, i.e., it uses the *n* closest neighbors concerning the architectural similarity. The effect of this is the discovery of niches that are not efficient in terms of fitness. In the scenario of global competition, individuals have to outperform a broader range of solutions in order to survive through generations, leading to the convergence of better models.

Figures 4 and 5 display the average number of samples used in the training process for all discriminators and generators in the population, respectively. In these charts, we confirm the effect of the novelty strategy applied in our solution, which is more evident in Figure 4. These results evidence that newer individuals were more frequently selected through generations in the solutions based on the QD algorithm, resulting in fewer training samples directly seen by them (new individuals can have new genes introduced by variation operators). As expected, we also show that novelty is more present in the local competition solution when compared to the global competition version.

Figures 6 and 7 provide additional support for this analysis, showing the number of samples used in training of the best individuals in the population of discriminators and generators, respectively. The curves follow a similar behavior



Figure 4: Average number of samples used to train all discriminators with MNIST.



Figure 5: Average number of samples used to train all generators with MNIST.



Figure 6: Average number of samples used to train best discriminators with MNIST.

of Figures 4 and 5, evidencing that best individuals at each generation also present more innovation in the solutions using COEGAN+NSLC.

A difference in the novelty effect is evident when comparing Figure 4 to Figure 5 and Figure 6 to Figure 7. The effect of innovation is more evident for discriminators (Figures 4 and 6) than in the results with generators (Figures 5 and 7). We



Figure 7: Average number of samples used to train best generators with MNIST.

attribute this difference to the choice of fitness functions. As concluded in [4], the FID score used in generators is a more reliable metric than the loss function used in discriminators. This affects the quality criterion used in the NSGA-II optimization method, making the selection of better individuals more assertive. However, further experiments are required to confirm this effect on innovation in the populations.

To better study the quality and diversity achieved by our model, we present in Figure 8 the distribution of samples produced by the best generator at different steps of the training process in one execution using the COEGAN+NSGC approach. Samples are placed in a  $40 \times 40$  grid, positioned according to t-SNE<sup>2</sup>. For this, 1600 samples from each case were used in the t-SNE training and a discretization function is applied to place these samples into the two-dimensional space. This method results in some overlapping samples, which indicates the level of diversity obtained by a model, i.e., fewer overlapping samples is evidence of better representation of the latent space.

Figure 8(a) represents the distribution of the MNIST dataset, i.e., the variety of samples used in training. We can see in Figure 8(b) that, at the initial stage, the samples are noisy and do not resemble images from MNIST, creating a high number of overlapping samples. In generation 10, represented by Figure 8(c), the distribution of samples is more close to the presented in Figure 8(a), although we can still see some lack of quality and under-representation of some digits. Figure 8(d) shows samples produced after the whole evolutionary process. These samples have better quality and preserve diversity, resulting in 1077 overlapping samples, even lower than the 1121 overlapping samples presenting in the MNIST dataset.

To assess the efficiency of our solution in complex datasets, we used COEGAN+NSGC, the best performing version of the proposed algorithm, to conduct experiments with CelebA [15]. For the sake of simplicity, we reduced the type of layers only to convolution and transpose convolution layers when adding GECCO '20, July 8-12, 2020, Cancún, Mexico



1189 overlapped samples

(d) Generation 501077 overlapped samples

Figure 8: Distribution of samples using t-SNE with the MNIST Dataset. We show samples (a) from the input dataset, the best generator at the (b) first generation, (c) after ten generations, and (d) at the end of training. We fed t-SNE with 1600 samples from each scenario and used the results for positioning them into a two-dimensional space. The number of overlapped samples is displayed for each case.



Figure 9: Best FID Score on the CelebA dataset.

a new gene, excluding the linear layer from the set of possibilities. Besides, the activation functions were restricted to ReLU and Leaky ReLU in the mutation operators. The populations of generators and discriminators contain 5 individuals each.

 $<sup>^2</sup>An$  expanded version of these images using a  $120\times120$  grid is available at https://github.com/vfcosta/qd-coegan.

The images from the CelebA dataset were rescaled to  $64 \times 64$ . To handle images of bigger sizes, we increased the genome limit to 5 and the number of batches per generation to 200. The remainder of the parameters is the same presented in Table 1.

We compare the results of our approach with a regular GAN that uses an architecture based on DCGAN [22]. In the DCGAN-based experiments, the architecture of the generator and the discriminator is composed of four layers. Previous experiments with neural networks using five layers were conducted but the results were more unstable, making the four-layers version more suitable for comparison with CO-EGAN+NSGC. It is important to note that we ensure the DCGAN approach is trained with the same number of samples of an individual in COEGAN+NSGC. Therefore, we define one training epoch as the training of DCGAN with 1000 batches. For COEGAN+NSGC, one training epoch is equivalent to one generation of the evolutionary algorithm. As we use the *all vs. all* pairing approach, each individual is also trained with 1000 batches per generation (5 individuals times 200 batches).

Figure 9 presents the progression of the best FID score when training with CelebA. We can see a smooth progression of the FID in the COEGAN+NSGC approach, leading to a final result consistently better when compared to the DCGAN-based solution. In DCGAN, FID varies during the training epoch, demonstrating spikes during the process mostly due to the occurrence of common stability issues on the GAN training, such as the mode collapse problem [3]. This is evidence that our approach provides more stability on GAN training when compared to regular GANs with similar architectures.

Figure 10 shows samples created by both approaches when trained with the CelebA dataset. Figure 10(a) displays samples created by the DCGAN approach after the final epoch when issues were observed in training. This is an example of the resulting effect of a spike that occurred in the DCGAN training. In Figure 10(b) we can see samples produced after training with COEGAN+NSGC. Although spikes are not present, the quality of samples produced by COEGAN+NSGC is not perfect, being worse than stateof-the-art GANs trained with CelebA [29]. Further experiments should be executed to assess the capability of CO-EGAN+NSGC to achieve better results concerning the FID score in larger experimental setups.

#### 5 CONCLUSIONS

In this paper, we have investigated the application of a quality diversity algorithm to train and evolve Generative Adversarial Networks (GANs). GANs are capable of producing strong generative models, but the training procedure is hard and stability issues frequently affect the results. Furthermore, the hand-design of efficient models is a time-consuming task, worsen by the inconstancy of the GAN training.

Evolutionary algorithms were recently proposed to improve the training of GANs and to provide the discovery of



Figure 10: Samples created by (a) DCGAN after collapsing in final training epoch and by (b) CO-EGAN+NSGC after training.

efficient generative models. COEGAN uses coevolution with an evolutionary algorithm inspired by NEAT to train GANs. However, the lack of diversity and premature optimization leave room for improvement of the solution.

We propose in this paper the extension of COEGAN to use a Quality Diversity (QD) algorithm in order to improve the exploration of the search space. Therefore, we design a new evolutionary algorithm that combines COEGAN with the approach used in the Novelty Search with Local Competition (NSLC) algorithm.

The experimental results show that the use of QD to guide the evolution of GANs improved the diversity in the population, leading to the discovery of better models. Furthermore, experiments with a version of the algorithm using global competition evidence that we can consistently outperform the previous results of COEGAN in the MNIST dataset. Experiments with the CelebA dataset indicate that our proposal provides a more stable training when compared to a regular GAN based on the DCGAN architecture, avoiding problems such as mode collapse and vanishing gradient.

As further works, we pretend to explore our quality diversity approach and extend the experimental setup to increase the capability to discover better models when training in complex datasets. Besides, new architectural components recently proposed for GANs can be incorporated into the model to enhance the population of individuals represented by our solution.

### ACKNOWLEDGMENTS

This work is funded by national funds through the FCT - Foundation for Science and Technology, I.P., within the scope of the project CISUC - UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020.

#### REFERENCES

- Abdullah Al-Dujaili, Tom Schmiedlechner, Erik Hemberg, and Una-May O'Reilly. 2018. Towards distributed coevolutionary GANs. In AAAI 2018 Fall Symposium.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In International

 $Conference\ on\ Machine\ Learning.\ 214-223.$ 

- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019.
- [4] Victor Costa, Nuno Lourenço, João Correia, and Penousal Machado. 2019. COEGAN: Evaluating the coevolution effect in generative adversarial networks. In *Proceedings of the Genetic* and Evolutionary Computation Conference. ACM, 374–382.
- [5] Victor Costa, Nuno Lourenço, and Penousal Machado. 2019. Coevolution of Generative Adversarial Networks. In International Conference on the Applications of Evolutionary Computation (Part of EvoStar). Springer, 473–487.
- [6] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [7] William Fedus, Mihaela Rosca, Balaji Lakshminarayanan, Andrew M Dai, Shakir Mohamed, and Ian Goodfellow. 2018. Many paths to equilibrium: GANs do not need to decrease a divergence at every step. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 -May 3, 2018, Conference Track Proceedings.
- [8] Unai Garciarena, Roberto Santana, and Alexander Mendiburu. 2018. Evolved GANs for Generating Pareto Set Approximations. In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '18). ACM, New York, NY, USA, 434-441.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In Advances in Neural Information Processing Systems. 2672–2680.
- [10] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of Wasserstein GANs. In Advances in Neural Information Processing Systems. 5769–5779.
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In Advances in Neural Information Processing Systems. 6629–6640.
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In 6th International Conference on Learning Representations, ICLR 2018, Vancower, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings.
- [13] Yann LeCun. 1998. The MNIST database of handwritten digits. http://yann.lecun.com/exdb/mnist/ (1998).
- [14] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation. 211–218.
- [15] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In Proceedings of the IEEE International Conference on Computer Vision. 3730– 3738.

- [16] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. Journal of Machine Learning Research 9, Nov (2008), 2579–2605.
- [17] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least squares generative adversarial networks. In 2017 IEEE International Conference on Computer Vision (ICCV). IEEE, 2813-2821.
- [18] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Dan Fink, Olivier Francon, Bala Raju, Arshak Navruzyan, Nigel Duffy, and Babak Hodjat. 2017. Evolving Deep Neural Networks. arXiv preprint arXiv:1703.00548 (2017).
- [19] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral Normalization for Generative Adversarial Networks. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 -May 3, 2018, Conference Track Proceedings.
- [20] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. arXiv preprint arXiv:1504.04909 (2015).
- [21] Justin K Pugh, Lisa B Soros, Paul A Szerlip, and Kenneth O Stanley. 2015. Confronting the challenge of quality diversity. In Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation. 967–974.
- [22] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015).
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [24] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training GANs. In Advances in Neural Information Processing Systems. 2234-2242.
- [25] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary Compu*tation 10, 2 (2002), 99–127.
- [26] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- [27] Jamal Toutouh, Erik Hemberg, and Una-May O'Reilly. 2019. Spatial Evolutionary Generative Adversarial Networks. arXiv preprint arXiv:1905.12702 (2019).
- [28] Chaoyue Wang, Chang Xu, Xin Yao, and Dacheng Tao. 2018. Evolutionary Generative Adversarial Networks. arXiv preprint arXiv:1803.00657 (2018).
- [29] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. 2018. Stack-GAN++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (2018), 1947–1962.