

Evolutionary Generative Adversarial Networks

Chaoyue Wang, Chang Xu, Xin Yao, *Fellow, IEEE*, Dacheng Tao, *Fellow, IEEE*

Abstract—Generative adversarial networks (GAN) have been effective for learning generative models for real-world data. However, accompanied with the generative tasks becoming more and more challenging, existing GANs (GAN and its variants) tend to suffer from different training problems such as instability and mode collapse. In this paper, we propose a novel GAN framework called evolutionary generative adversarial networks (E-GAN) for stable GAN training and improved generative performance. Unlike existing GANs, which employ a pre-defined adversarial objective function alternately training a generator and a discriminator, we evolve a population of generators to play the adversarial game with the discriminator. Different adversarial training objectives are employed as mutation operations and each individual (*i.e.*, generator candidature) are updated based on these mutations. Then, we devise an evaluation mechanism to measure the quality and diversity of generated samples, such that only well-performing generator(s) are preserved and used for further training. In this way, E-GAN overcomes the limitations of an individual adversarial training objective and always preserves the well-performing offspring, contributing to progress in and the success of GANs. Experiments on several datasets demonstrate that E-GAN achieves convincing generative performance and reduces the training problems inherent in existing GANs.

Index Terms—Generative adversarial networks, evolutionary computation, deep generative models

I. INTRODUCTION

GENERATIVE adversarial networks (GAN) [1] are one of the main groups of methods used to learn generative models from complicated real-world data. As well as using a generator to synthesize semantically meaningful data from standard signal distributions, GANs (GAN and its variants) train a discriminator to distinguish *real* samples in the training dataset from *fake* samples synthesized by the generator. As the confronter, the generator aims to deceive the discriminator by producing ever more realistic samples. The training procedure continues until the generator wins the adversarial game; that is, the discriminator cannot make a better decision than randomly

guessing whether a particular sample is fake or real. Compared with other existing generative models, GANs provide a concise and efficient framework for learning generative models. Therefore, GANs have recently been successfully applied to image generation [2]–[6], image editing [7]–[10], video prediction [11]–[13], and many other tasks [14]–[18].

Comparing with most existing discriminative tasks (*e.g.*, classification, clustering), GANs perform a challenging generative process, which projects a standard distribution to a much more complex high-dimensional real-world data distribution. Firstly, although GANs have been utilized to modeling large-scale real-world datasets, such as CelebA, LSUN and ImageNet, they are easily suffering from mode collapse problem, *i.e.*, the generator can only learn some limited patterns from the large-scale target datasets, or assigns all its probability mass to a small region in the space [19]. Meanwhile, if the generated distribution and the target data distribution do not substantially overlap (usually at the beginning of training), the generator gradients can point to more or less random directions or even result in the vanishing gradient issue. In addition, to generate high-resolution and high-quality samples, both the generator and discriminator are asked to be deeper and larger. Under the vulnerable adversarial framework, it's hard to balance and optimize such large-scale deep networks. Thus, in most existing works, appropriate hyper-parameters (*e.g.*, learning rate and updating steps) and network architectures are critical configurations. Unsuitable settings reduce GAN's performance or even fail to produce any reasonable results. Overall, although GANs already produce visually appealing samples in various applications, they are still facing many large-scale optimization problems.

Many recent efforts on GANs have focused on overcoming these optimization difficulties by developing various adversarial training objectives. Typically, assuming the optimal discriminator for the given generator is learned, different objective functions of the generator aim to measure the distance between the generated distribution and the target data distribution under different metrics. The original GAN uses Jensen-Shannon divergence (JSD) as the metric. A number of metrics have been introduced to improve GAN's performance, such as least-squares [20], absolute deviation [21], Kullback-Leibler (KL) divergence [22], and Wasserstein distance [23]. However, according to both theoretical analyses and experimental results, minimizing each distance has its own pros and cons. For example, although measuring KL divergence largely eliminates the vanishing gradient issue, it easily results in mode collapse [22], [24]. Likewise, Wasserstein distance greatly improves training stability but can have non-convergent limit cycles near equilibrium [25].

Through observation, we find that most existing GAN methods are limited by the specified adversarial optimization

This work was supported by the Australian Research Council Projects: FL-170100117, DP-180103424, IH180100002, and DE180101438; and National Key R&D Program of China (Grant No. 2017YFC0804003), EPSRC (Grant Nos. EP/J017515/1 and EP/P005578/1), the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (Grant No. 2017ZT07X386), Shenzhen Peacock Plan (Grant No. KQTD2016112514355531), the Science and Technology Innovation Committee Foundation of Shenzhen (Grant No. ZDSYS201703031748284) and the Program for University Key Laboratory of Guangdong Province (Grant No. 2017KSYS008).

C. Wang, C. Xu and D. Tao are with the UBTECH Sydney Artificial Intelligence Centre and the School of Computer Science, in the Faculty of Engineering and Information Technologies, at the University of Sydney, 6 Cleveland St, Darlingtown, NSW 2008, Australia (email: chaoyue.wang, c.xu, dacheng.tao@sydney.edu.au).

X. Yao is with the Shenzhen Key Laboratory of Computational Intelligence, University Key Laboratory of Evolving Intelligent Systems of Guangdong Province, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China, and CERCA, School of Computer Science, University of Birmingham, UK (e-mail: x.yao@cs.bham.ac.uk).

strategy. Since the training strategy is fixed, it is hard to adjust the balance between the generator and discriminator during the training process. Meanwhile, as aforementioned, each existing adversarial training strategy has its own pros and cons in training GAN models.

In this paper, we build an evolutionary generative adversarial network (E-GAN), which treats the adversarial training procedure as an evolutionary problem. Specifically, a discriminator acts as the *environment* (i.e., provides adaptive loss functions) and a *population* of generators evolve in response to the environment. During each adversarial (or evolutionary) iteration, the discriminator is still trained to recognize real and fake samples. However, in our method, acting as parents, generators undergo different *mutations* to produce offspring to adapt to the environment. Different adversarial objective functions aim to minimize different distances between the generated distribution and the data distribution, leading to different mutations. Meanwhile, given the current optimal discriminator, we measure the quality and diversity of samples generated by the updated offspring. Finally, according to the principle of “survival of the fittest”, poorly-performing offspring are removed and the remaining well-performing offspring (i.e., generators) are preserved and used for further training. Based on the evolutionary paradigm to optimize GANs, the proposed E-GAN overcomes the inherent limitations in the individual adversarial training objectives and always preserves the well-performing offspring produced by different training objectives (i.e., mutations). In this way, we contribute to progress in and the success of the large-scale optimization of GANs. Following vanilla GAN [1], we evaluate the new algorithm in image generation tasks. Overall, the proposed evolutionary strategy is largely orthogonal to existing GAN models. Through applying the evolutionary framework to different GAN models, it is possible to perform different kinds of generation tasks. For example, GAN objectives were devised for generating text. Considering them as mutation operations, the proposed evolutionary framework can be applied to solve text generation tasks.

Meanwhile, the proposed E-GAN framework also provides a novel direction to apply evolutionary learning paradigm on solving deep learning problems. Recent years, although deep learning algorithms have achieved promising performance on a variety of applications, they are still facing many challenges in solving real-world problems. Evolutionary computation, as a powerful approach to complex real-world problems [26]–[29], has been utilized to solve many deep learning challenges. Among them, [30] devised an evolutionary algorithm to automatically search the architecture and hyper-parameters of deep networks. Moreover, evolution strategies have been utilized as an alternative to MDP-based techniques to optimize reinforcement learning models [31]. In this work, we attempt to combine the back-propagation algorithm and the evolutionary algorithm for optimizing deep generative models. The parameters updated by different learning objectives are regarded as variation results during the evolutionary process. By introducing suitable evaluation and selection mechanisms, the whole training process can be more efficiency and stability. We hope the proposed evolutionary learning framework can

be generalized to more deep learning problems, such as reinforcement learning.

In summary, we make following contributions in this paper:

- In order to stabilize GAN’s training process, we devised a simple yet efficient evolutionary algorithm for optimizing generators within GANs framework. To the best of our knowledge, it’s the first work that introduces the evolutionary learning paradigm into learning GAN models.
- Through analyzing the training process of E-GAN, some properties of existing GANs objectives are further explored and discussed.
- Experiments evaluated on several large-scale datasets are performed, and demonstrate that convincing results can be achieved by the proposed E-GAN framework.

The rest of the paper is organized as follows: after a brief summary of previous related works in section II, we illustrate the proposed E-GAN together with its training process in section III. Then we exhibit the experimental validation of the whole method in section IV. Finally, we conclude this paper with some future directions in section V.

II. BACKGROUND

In this section, we first review some previous GANs devoted to reducing training instability and improving the generative performance. We then briefly summarize some evolutionary algorithms on deep neural networks.

A. Generative Adversarial Networks

Generative adversarial networks (GAN) provides an excellent framework for learning deep generative models, which aim to capture probability distributions over the given data. Compared to other generative models, GAN is easily trained by alternately updating a generator and a discriminator using the back-propagation algorithm. In many generative tasks, GANs (GAN and its variants) produce better samples than other generative models [32]. Besides image generation tasks, GANs have been introduced to more and more tasks, such as video generation [11], [33], visual tracking [34]–[36], domain adaption [37], hashing coding [38]–[40], and feature learning [41], [42]. In these tasks, the adversarial training strategy also achieved promising performance.

However, some problems still exist in the GANs training process. In the original GAN, training the generator was equal to minimizing the Jensen-Shannon divergence between the data distribution and the generated distribution, which easily resulted in the vanishing gradient problem. To solve this issue, a non-saturating heuristic objective (i.e., ‘ $-\log D$ trick’) replaced the minimax objective function to penalize the generator [1]. Then, [22] and [43] designed specified network architectures (DCGAN) and proposed several heuristic tricks (e.g., feature matching, one-side label smoothing, virtual batch normalization) to improve training stability. Meanwhile, energy-based GAN [21] and least-squares GAN [20] improved training stability by employing different training objectives. Although these methods partly enhanced training stability, in practice, the network architectures and training procedure still required careful design to maintain the discriminator-generator

balance. More recently, Wasserstein GAN (WGAN) [23] and its variant WGAN-GP [44] were proposed to minimize the Wasserstein-1 distance between the generated and data distributions. Since the Wasserstein-1 distance is continuous everywhere and differentiable almost everywhere under only minimal assumptions [23], these two methods convincingly reduce training instability. However, to measure the Wasserstein-1 distance between the generated distribution and the data distribution, they are asked to enforce the Lipschitz constraint on the discriminator (*aka* critic), which may result in some optimization problems [44].

Besides devising different objective functions, some works attempt to stable GAN training and improve generative performance by introducing multiple generators or discriminators into an adversarial training framework. [45] proposed a dual discriminator GAN (D2GAN), which combines the Kullback-Leibler (KL) and reverse KL divergences into a unified objective function through employing two discriminators. Multi-discriminator GAN frameworks [46], [47] are devised and utilized for providing stable gradients to the generator and further stabilizing the adversarial training process. Moreover, [48] apply boosting techniques to train a mixture of generators by continually adding new generators to the mixture. [49] train many generators by using a multi-class discriminator that predicts which generator produces the sample. Mixture GAN (MGAN) [50] is proposed to overcome the mode collapse problem by training multiple generators to specialize in different data modes. Overall, these multi-generator GANs aim to learn a set of generators, and the mixture of their learned distributions would approximate the data distribution, *i.e.*, different generators are encouraged to capture different data modes. Although the proposed E-GAN also creates multiple generators during training, we always keep the well-performing candidates through survival of the fittest, which helps the final learned generator achieving better performance. Note that, in our framework, only one generator was learned to represent the whole target distribution at the end of the training.

B. Evolutionary Computation

Over the last twenty years, evolutionary algorithms have achieved considerable success across a wide range of computational tasks including modeling, optimization and design [51]–[55]. Inspired by natural evolution, the essence of an evolutionary algorithm is to equate possible solutions to individuals in a population, produce offspring through variations, and select appropriate solutions according to fitness [56].

Recently, evolutionary algorithms have been introduced to solve deep learning problems. To minimize human participation in designing deep algorithms and automatically discover such configurations, there have been many attempts to optimize deep learning hyper-parameters and design deep network architectures through an evolutionary search [57], [58]. Among them, [30] proposed a large-scale evolutionary algorithm to design a whole deep classifier automatically. Meanwhile, different from widely employed gradient-based learning algorithms (*e.g.*, backpropagation), evolutionary algorithms have

also demonstrated their capacity to optimize neural networks. EPNet [59] was devised for evolving and training neural networks using evolutionary programming. In [60], EvoAE was proposed to speed up the training of autoencoders for constructing deep neural networks. Moreover, [31] proposed a novel evolutionary strategy as an alternative to the popular MDP-based reinforcement learning techniques, achieving strong performance on reinforcement learning benchmarks. In addition, an evolutionary algorithm was proposed to compress deep learning models by automatically eliminating redundant convolution filters [61]. Last but not the least, the evolutionary learning paradigm has been utilized to solve a number of deep/machine tasks, such as automatic machine learning [62], multi-objective optimization [63], *etc.* However, to best of our knowledge, there is still no work attempt to optimize deep generative models with evolutionary learning algorithms.

III. METHODS

In this section, we first review the original GAN formulation. Then, we introduce the proposed E-GAN algorithm. By illustrating E-GAN's mutations and evaluation mechanism, we further discuss the advantage of the proposed framework. Finally, we conclude with the entire E-GAN training process.

A. Generative Adversarial Networks

GAN, first proposed in [1], studies a two-player minimax game between a discriminative network D and a generative network G . Taking noisy sample $z \sim p(z)$ (sampled from a uniform or normal distribution) as the input, the generative network G outputs new data $G(z)$, whose distribution p_g is supposed to be close to that of the data distribution p_{data} . Meanwhile, the discriminative network D is employed to distinguish the true data sample $x \sim p_{\text{data}}(x)$ and the generated sample $G(z) \sim p_g(G(z))$. In the original GAN, this adversarial training process was formulated as:

$$\min_G \max_D \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))]. \quad (1)$$

The adversarial procedure is illustrated in Fig. 1 (a). Most existing GANs perform a similar adversarial procedure in different adversarial objective functions.

B. Evolutionary Algorithm

In contrast to conventional GANs, which alternately update a generator and a discriminator, we devise an evolutionary algorithm that evolves a population of generator(s) $\{G\}$ in a given environment (*i.e.*, the discriminator D). In this population, each *individual* represents a possible solution in the parameter space of the generative network G . During the evolutionary process, we expect that the population gradually adapts to its environment, which means that the evolved generator(s) can generate ever more realistic samples and eventually learn the real-world data distribution. As shown in Fig. 1 (b), during evolution, each step consists of three sub-stages:

- **Variation:** Given an individual G_θ in the population, we utilize the variation operators to produce its offspring $\{G_{\theta_1}, G_{\theta_2}, \dots\}$. Specifically, several copies of

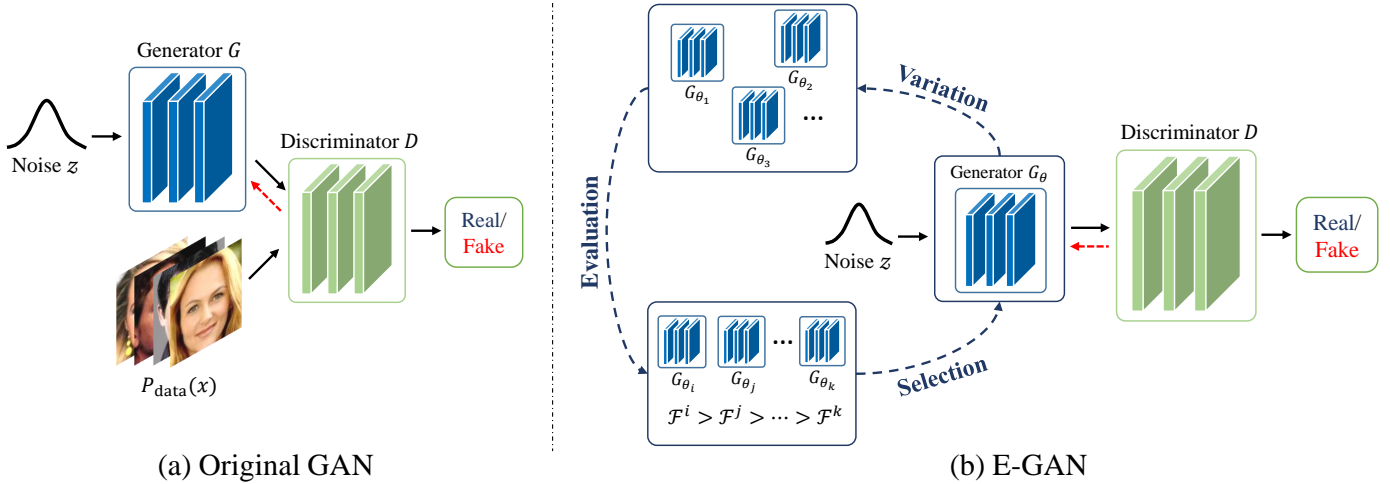


Fig. 1. (a) The original GAN framework. A generator G and a discriminator D play a two-player adversarial game. The updating gradients of the generator G are received from the adaptive objective, which depends on discriminator D . (b) The proposed E-GAN framework. A population of generators $\{G_{\theta}\}$ evolves in a dynamic environment, the discriminator D . Each evolutionary step consists of three sub-stages: variation, evaluation, and selection. The best offspring are kept.

each individual—or *parent*—are created, each of which are modified by different *mutations*. Then, each modified copy is regarded as one *child*.

- **Evaluation:** For each child, its performance—or *individual's quality*—is evaluated by a *fitness function* $\mathcal{F}(\cdot)$ that depends on the current environment (*i.e.*, discriminator D).
- **Selection:** All children will be selected according to their fitness value, and the worst part is removed—that is, they are *killed*. The rest remain *alive* (*i.e.*, free to act as parents), and evolve to the next iteration.

After each evolutionary step, the discriminative network D (*i.e.*, the environment) is updated to further distinguish real samples x and fake samples y generated by the evolved generator(s), *i.e.*,

$$\mathcal{L}_D = -\mathbb{E}_{x \sim p_{data}} [\log D(x)] - \mathbb{E}_{y \sim p_g} [\log(1 - D(y))]. \quad (2)$$

Thus, the discriminative network D (*i.e.*, the environment) can continually provide the adaptive losses to drive the population of generator(s) evolving to produce better solutions. Next, we illustrate and discuss the proposed variation (or mutation), evaluation and selection operators in detail.

C. Variation

We employ *asexual reproduction* with different mutations to produce the next generation's individuals (*i.e.*, children). Specifically, these mutation operators correspond to different training objectives, which attempt to narrow the distances between the generated distribution and the data distribution from different perspectives. In this section, we introduce the muta-

tions used in this work¹. To analyze the corresponding properties of these mutations, we suppose that, for each evolutionary step, the optimal discriminator $D^*(x) = \frac{p_{data}(x)}{p_{data}(x) + p_g(x)}$, according to Eq. (2), has already been learned [1].

1) *Minimax mutation:* The minimax mutation corresponds to the minimax objective function in the original GAN:

$$\mathcal{M}_G^{\minimax} = \frac{1}{2} \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))]. \quad (3)$$

According to the theoretical analysis in [1], given the optimal discriminator D^* , the minimax mutation aims to minimize the Jensen-Shannon divergence (JSD) between the data distribution and the generated distribution. Although the minimax game is easy to explain and theoretically analyze, its performance in practice is disappointing, a primary problem being the generator's vanishing gradient. If the support of two distributions lies in two manifolds, the JSD will be a constant, leading to the vanishing gradient [24]. This problem is also illustrated in Fig. 2. When the discriminator rejects generated samples with high confidence (*i.e.*, $D(G(z)) \rightarrow 0$), the gradient tends to vanishing. However, if the generated distribution overlaps with the data distribution, meaning that the discriminator cannot completely distinguish real from fake samples, the minimax mutation provides effective gradients and continually narrows the gap between the data distribution and the generated distribution.

2) *Heuristic mutation:* Unlike the minimax mutation, which minimizes the log probability of the discriminator being

¹Although more mutation operations can be included in our framework, according to the theoretical analysis below, we adopt three interpretable and complementary objectives as our mutations. Meanwhile, we have tested more mutation operations, yet the mutations described in this paper already delivered a convincing performance.




Fig. 2. The mutation (or objective) functions that the generator G receives given the discriminator D .

correct, the heuristic mutation aims to maximize the log probability of the discriminator being mistaken, *i.e.*,

$$\mathcal{M}_G^{\text{heuristic}} = -\frac{1}{2} \mathbb{E}_{z \sim p_z} [\log(D(G(z)))]. \quad (4)$$

Compared to the minimax mutation, the heuristic mutation will not saturate when the discriminator rejects the generated samples. Thus, the heuristic mutation avoids vanishing gradient and provides useful generator updates (Fig. 2). However, according to [24], given the optimal discriminator D^* , minimizing the heuristic mutation is equal to minimizing $[KL(p_g || p_{\text{data}}) - 2JSD(p_g || p_{\text{data}})]$, *i.e.*, inverted KL minus two JSDs. Intuitively, the JSD sign is negative, which means pushing these two distributions away from each other. In practice, this may lead to training instability and generative quality fluctuations [44].

3) *Least-squares mutation*: The least-squares mutation is inspired by LSGAN [20], where the least-squares objectives are utilized to penalize its generator to deceive the discriminator. In this work, we formulate the least-squares mutation as:

$$\mathcal{M}_G^{\text{least-square}} = \mathbb{E}_{z \sim p_z} [(D(G(z)) - 1)^2]. \quad (5)$$

As shown in Fig. 2, the least-squares mutation is non-saturating when the discriminator can recognize the generated sample (*i.e.*, $D(G(z)) \rightarrow 0$). When the discriminator output grows, the least-squares mutation saturates, eventually approaching zero. Therefore, similar to the heuristic mutation, the least-squares mutation can avoid vanishing gradient when the discriminator has a significant advantage over the generator. Meanwhile, compared to the heuristic mutation, although the least-squares mutation will not assign an extremely high cost to generate fake samples, it will also not assign an extremely low cost to mode dropping², which partly avoids mode collapse [20].

² [24] demonstrated that the heuristic objective suffers from mode collapse since $KL(p_g || p_{\text{data}})$ assigns a high cost to generating fake samples but an extremely low cost to mode dropping.

Note that, different from GAN-minimax and GAN-heuristic, LSGAN employs a different objective ('least-squares') to optimize the discriminator, *i.e.*,

$$\begin{aligned} \mathcal{L}_D^{\text{LSGAN}} &= \frac{1}{2} \mathbb{E}_{x \sim p_{\text{data}}} [(D(x) - 1)^2] + \frac{1}{2} \mathbb{E}_{z \sim p_z} [D(G(z))^2] \\ &= \int_x \frac{1}{2} (p_{\text{data}}(x)(D(x) - 1)^2 + p_g(x)D(x)^2) dx. \end{aligned} \quad (6)$$

Yet, with respect to $D(x)$, the function $\mathcal{L}_D^{\text{LSGAN}}$ achieves its minimum in $[0, 1]$ at $\frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_g(x)}$, which is equivalent to ours (*i.e.*, Eq. (2)).

Therefore, although we employ only one discriminator as the environment to distinguish real and generated samples, it is sufficient to provide adaptive losses for all mutations described above.

D. Evaluation

In an evolutionary algorithm, evaluation is the operation of measuring the quality of individuals. To determine the evolutionary direction (*i.e.*, individuals' selection), we devise an evaluation (or fitness) function to measure the performance of evolved individuals (*i.e.*, children). Typically, we mainly focus on two properties: 1) the quality and 2) the diversity of generated samples. Quality is measured for each generated sample. If a generated image could be realistic enough, it will fool the discriminator. On the other hand, the diversity measures whether the generator could spread the generated samples out enough, which could largely avoid mode collapse.

Firstly, we simply feed generator produced images into the discriminator D and observe the average value of the output, which we name the *quality fitness score*:

$$\mathcal{F}_q = \mathbb{E}_z [D(G(z))]. \quad (7)$$

Note that discriminator D is constantly upgraded to be optimal during the training process, reflecting the quality of generators at each evolutionary (or adversarial) step. If a generator obtains a relatively high quality score, its generated samples can deceive the discriminator and the generated distribution is further approximate to the data distribution.

Besides generative quality, we also pay attention to the diversity of generated samples and attempt to overcome the mode collapse issue in GAN optimization. Recently, [25] proposed a gradient-based regularization term to stabilize the GAN optimization and suppress mode collapse. When the generator collapses to a small region, the discriminator will subsequently label collapsed points as fake with obvious countermeasure (*i.e.*, big gradients). In contrast, if the generator is capable of spreading generated data out enough, the discriminator will not be much confident to label generated samples as fake data (*i.e.*, updated with small gradients). Other techniques, *e.g.* exploiting conditioning latent vector [64] or subspace [65], can also be applied to pursue diversity. But this issue does not fall within the scope of this paper.

We employ a similar principle to evaluate generator optimization stability and generative diversity. Here, since the gradient-norm of the discriminator could vary largely during training, we employed logarithm to shrink its fluctuation.

Specifically, the minus log-gradient-norm of optimizing D is utilized to measure the diversity of generated samples. If an evolved generator obtains a relatively high value, which corresponds to small discriminator gradients, its generated samples tend to spread out enough, to avoid the discriminator from having obvious countermeasures. Thus, the mode collapse issue can be suppressed and the discriminator will change smoothly, which helps to improve the training stability. Formally, the *diversity fitness score* is defined as:

$$\mathcal{F}_d = -\log \left| \left| \nabla_D - \mathbb{E}_x[\log D(x)] - \mathbb{E}_z[\log(1 - D(G(z)))] \right| \right|. \quad (8)$$

Based on the aforementioned two fitness scores, we can finally give the evaluation (or fitness) function of the proposed evolutionary algorithm:

$$\mathcal{F} = \mathcal{F}_q + \gamma \mathcal{F}_d, \quad (9)$$

where $\gamma \geq 0$ balances two measurements: generative quality and diversity. Overall, a relatively high fitness score \mathcal{F} , leads to higher training efficiency and better generative performance.

It is interesting to note that, the proposed fitness score can also be regarded as an objective function for generator G . However, as demonstrated in our experiments, without the proposed evolutionary strategy, this fitness objective is still hard to achieve convincing performance. Meanwhile, similar with most existing objective function of GANs, the fitness function will continuously fluctuate during the dynamic adversarial training process. Specifically, in each iteration, the fitness score of generators are evaluated based on the current discriminator. Comparing them, we are capable of selecting well-performing ones. However, since the discriminator will be updated according to different generators, it is hard to discuss the connection between the fitness scores over different iterations.

E. Selection

In an evolutionary algorithm, the counterpart of the mutation operators is the selection. In the proposed E-GAN, we employ a simple yet useful survivor selection strategy to determine the next generation based on the fitness score of existing individuals.

First of all, we should notice that both the generators (*i.e.*, population) and the discriminator (*i.e.*, environment) are optimized alternately in a dynamic procedure. Thus, the fitness function is not fixed and the fitness score of generators can be only evaluated by the corresponding discriminator in the same evolutionary generation, which means fitness scores that evaluated in different generations cannot compare with each other. In addition, due to the mutation operators of the proposed E-GAN actually correspond to different adversarial training objectives, selecting desired offspring is equivalent to selecting the effective adversarial strategies. During the adversarial process, we hope that generator(s) can do it best to fool the discriminator (*i.e.*, implement the optimal adversarial strategy). Considered these two points, we utilize the comma selection, *i.e.*, (μ, λ) -selection [66] as the selection mechanism of E-GAN. Specifically, after sorting the current offspring population $\{\mathbf{x}_i\}_{i=1}^\lambda$ according to their fitness scores

Algorithm 1 Evolutionary generative adversarial networks (E-GAN). Default values $\alpha = 0.0002$, $\beta_1 = 0.5$, $\beta_2 = 0.99$, $n_D = 3$, $n_m = 3$, $m = 32$.

Require: the batch size m . the discriminator's updating steps per iteration n_D . the number of parents μ . the number of mutations n_m . Adam hyper-parameters α, β_1, β_2 , the hyper-parameter γ of evaluation function.

Require: initial discriminator's parameters w_0 . initial generators' parameters $\{\theta_0^1, \theta_0^2, \dots, \theta_0^\mu\}$.

```

1: for number of training iterations do
2:   for  $k = 0, \dots, n_D$  do
3:     Sample a batch of  $\{x^{(i)}\}_{i=1}^m \sim p_{\text{data}}$  (training data),
       and a batch of  $\{z^{(i)}\}_{i=1}^m \sim p_z$  (noise samples).
4:      $g_w \leftarrow \nabla_w [\frac{1}{m} \sum_{i=1}^m \log D_w(x^{(i)})$ 
5:            $+ \frac{1}{m} \sum_{j=1}^\mu \sum_{i=1}^m \log(1 - D_w(G_{\theta^j}(z^{(i)})))]$ 
6:      $w \leftarrow \text{Adam}(g_w, w, \alpha, \beta_1, \beta_2)$ 
7:   end for
8:   for  $j = 0, \dots, \mu$  do
9:     for  $h = 0, \dots, n_m$  do
10:      Sample a batch of  $\{z^{(i)}\}_{i=1}^m \sim p_z$ 
11:       $g_{\theta^j, h} \leftarrow \nabla_{\theta^j} \mathcal{M}_G^h(\{z^{(i)}\}_{i=1}^m, \theta^j)$ 
12:       $\theta_{\text{child}}^{j, h} \leftarrow \text{Adam}(g_{\theta^j, h}, \theta^j, \alpha, \beta_1, \beta_2)$ 
13:       $\mathcal{F}^{j, h} \leftarrow \mathcal{F}_q^{j, h} + \gamma \mathcal{F}_d^{j, h}$ 
14:    end for
15:  end for
16:   $\{\mathcal{F}^{j_1, h_1}, \mathcal{F}^{j_2, h_2}, \dots\} \leftarrow \text{sort}(\{\mathcal{F}^{j, h}\})$ 
17:   $\theta^1, \theta^2, \dots, \theta^\mu \leftarrow \theta_{\text{child}}^{j_1, h_1}, \theta_{\text{child}}^{j_2, h_2}, \dots, \theta_{\text{child}}^{j_\mu, h_\mu}$ 
18: end for
```

\mathcal{F}_i , the μ -best individuals are selected to form the next generation.

F. E-GAN

Having introduced the proposed evolutionary algorithm and corresponding mutation operations, evaluation criteria and selection strategy, the complete E-GAN training process is concluded in Algorithm 1. Overall, in E-GAN, generators $\{G\}$ are regarded as an evolutionary population and discriminator D acts as an environment. For each evolutionary step, generators are updated with different mutations (or objectives) to accommodate the current environment. According to the principle of “survival of the fittest”, only well-performing children will survive and participate in future adversarial training. Unlike the two-player game with a fixed and static adversarial training objective in conventional GANs, E-GAN allows the algorithm to integrate the merits of different adversarial objectives and generate the most competitive solution. Thus, during training, the evolutionary algorithm not only largely suppresses the limitations (vanishing gradient, mode collapse, *etc.*) of individual adversarial objectives, but it also harnesses their advantages to search for a better solution.

IV. EXPERIMENTS

To evaluate the proposed E-GAN, we run experiments on several generative tasks and present the experimental results in this section. Compared with some previous GAN methods,

TABLE I
THE ARCHITECTURES OF THE GENERATIVE AND DISCRIMINATIVE NETWORKS.

Generative network G	
Input: Noise z , 100	
[layer 1]	Fully connect and Reshape to $(4 \times 4 \times 2048)$; $LReLU$;
[layer 2]	Transposed Conv. (4, 4, 2048), stride=2; $LReLU$;
[layer 3]	Transposed Conv. (4, 4, 1024), stride=1; $LReLU$;
[layer 4]	Transposed Conv. (4, 4, 1024), stride=2; $LReLU$;
[layer 5]	Transposed Conv. (4, 4, 512), stride=1; $LReLU$;
[layer 6]	Transposed Conv. (4, 4, 512), stride=2; $LReLU$;
[layer 7]	Transposed Conv. (4, 4, 256), stride=1; $LReLU$;
[layer 8]	Transposed Conv. (4, 4, 256), stride=2; $LReLU$;
[layer 9]	Transposed Conv. (4, 4, 128), stride=1; $LReLU$;
[layer 10]	Transposed Conv. (4, 4, 128), stride=2; $LReLU$;
[layer 11]	Transposed Conv. (3, 3, 3), stride=1; $Tanh$;
Output: Generated Image, $(128 \times 128 \times 3)$	
Discriminative network D	
Input: Image $(128 \times 128 \times 3)$	
[layer 1]	Conv. (3, 3, 128), stride=1; Batchnorm; $LReLU$;
[layer 2]	Conv. (3, 3, 128), stride=2; Batchnorm; $LReLU$;
[layer 3]	Conv. (3, 3, 256), stride=1; Batchnorm; $LReLU$;
[layer 4]	Conv. (3, 3, 256), stride=2; Batchnorm; $LReLU$;
[layer 5]	Conv. (3, 3, 512), stride=1; Batchnorm; $LReLU$;
[layer 6]	Conv. (3, 3, 512), stride=2; Batchnorm; $LReLU$;
[layer 7]	Conv. (3, 3, 1024), stride=1; Batchnorm; $LReLU$;
[layer 8]	Conv. (3, 3, 1024), stride=2; Batchnorm; $LReLU$;
[layer 9]	Conv. (3, 3, 2048), stride=1; Batchnorm; $LReLU$;
[layer 10]	Conv. (3, 3, 2048), stride=2; Batchnorm; $LReLU$;
[layer 11]	Fully connected (1); $Sigmoid$;
Output: Real or Fake (Probability)	

we show that the proposed E-GAN can achieve impressive generative performance on large-scale image datasets.

A. Implementation Details

We evaluate E-GAN on two synthetic datasets and three image datasets: CIFAR-10 [67], LSUN bedroom [68], and CelebA [69]. For fair comparisons, we adopted the same network architectures with existing works [22], [44]. In addition, to achieve better performance on generating 128×128 images, we slightly modified both the generative network and the discriminator network based on the DCGAN architecture. Specifically, the batch norm layers are removed from the generator, and more features channels are applied to each convolutional layers. The detailed networks are listed in Table I, note that the network architectures of the other comparison experiments can be easily found in the referenced works.

We use the default hyper-parameter values listed in Algorithm 1 for all experiments. Note that the hyper-parameter γ is utilized to balance the measurements of samples quality (*i.e.*, \mathcal{F}_q) and diversity (*i.e.*, \mathcal{F}_d). Usually, the quality fitness score

\mathcal{F}_q lies in $[0, 1]$, while the diversity fitness score \mathcal{F}_d measures the log-gradient-norm of the discriminator D , which can vary largely according to D 's scale. Therefore, we first determine γ 's range based on the selected discriminator D . Then, we run grid search to find its value. In practice, we choose $\gamma = 0.5$ for the synthetic datasets, and $\gamma = 0.001$ for real-world data. In addition, recently, some works [44], [70] proposed the gradient penalty (GP) term to regularize the discriminator to provide precise gradients for updating the generator. Within the adversarial training framework, our contributions are largely orthogonal to the GP term. In our experiments, through the setting without GP term, we demonstrated the efficiency of the proposed method. Then, after introducing the GP term, the generative performance was further improved, which demonstrated our framework could also benefit from the regularization technique for the discriminator. Furthermore, all experiments were trained on Nvidia Tesla V100 GPUs. To train a model for 64×64 images using the DCGAN architecture cost around 20 hours on a single GPU.

B. Evaluation metrics

Besides directly reported generated samples of the learned generative networks, we choose the Maximum Mean Discrepancy (MMD) [71], [72], the Inception score (IS) [43] and the Fréchet Inception distance (FID) [73] as quantitative metrics. Among them, the MMD can be utilized to measure the discrepancy between the generated distribution and the target distribution for synthetic Gaussian mixture datasets. However, the MMD is difficult to directly apply to high-dimensional image datasets. Therefore, through applying the pre-trained Inception v3 network [74] to generated images, the IS computes the KL divergence between the conditional class distribution and the marginal class distribution. Usually, this score correlates well with the human scoring of the realism of generated images from the CIFAR-10 dataset, and a higher value indicates better image quality. However, some recent works [6], [75] revealed serious limitations of the IS, *e.g.*, the target data distribution (*i.e.*, the training data) has not been considered. In our experiments, we utilize the IS to measure the E-GAN's performance on the CIFAR-10, and compare our results to previous works. Moreover, FID is a more principled and reliable metric and has demonstrated better correlations with human evaluation for other datasets. Specifically, FID calculates the Wasserstein-2 distance between the generated images and the real-world images in the high-level feature space of the pre-trained Inception v3 network. Note that lower FID means closer distances between the generated distribution and the real-world data distribution. In all experiments, we randomly generated 50k samples to calculate the MMD, IS and FID.

C. Synthetic Datasets and Mode Collapse

In the first experiment, we adopt the experimental design proposed in [76], which trains GANs on 2D Gaussian mixture distributions. The mode collapse issue can be accurately measured on these synthetic datasets, since we can clearly observe and measure the generated distribution and the target data

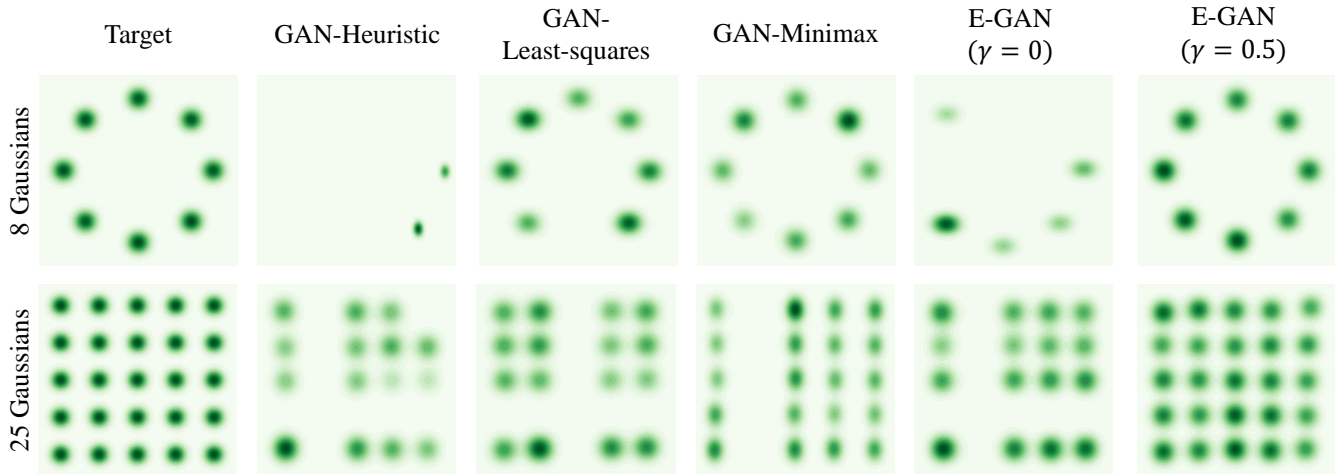


Fig. 3. Kernel density estimation (KDE) plots of the target data and generated data from different GANs trained on mixtures of Gaussians. In the first row, a mixture of 8 Gaussians arranged in a circle. In the second row, a mixture of 25 Gaussians arranged in a grid.

TABLE II

MMD ($\times 10^{-2}$) WITH MIXED GAUSSIAN DISTRIBUTIONS ON OUR TOY DATASETS. WE RAN EACH METHOD FOR 10 TIMES, AND REPORT THEIR AVERAGE AND BEST RESULTS. THE METHOD WITH LOWER MMD VALUE IMPLIES THE GENERATED DISTRIBUTION IS CLOSER TO THE TARGET ONE.

Methods	8 Gaussians		25 Gaussians	
	Average	Best	Average	Best
GAN-Heuristic	45.27	33.2	2.80	2.19
GAN-Least-squares	3.99	3.16	1.83	1.72
GAN-Minimax	2.94	1.89	1.65	1.55
E-GAN ($\lambda = 0$, without GP)	11.54	7.31	1.69	1.60
E-GAN ($\lambda = 0.5$, without GP)	2.36	1.17	1.20	1.04

distribution. As shown in Fig. 3, we employ two challenging distributions to evaluate E-GAN, a mixture of 8 Gaussians arranged in a circle and a mixture of 25 Gaussians arranged in a grid.³ Here, to evaluate if the proposed diversity fitness score can reduce the mode collapse, we did not introduce the gradient penalty norm and set the survived parents number μ as 1, *i.e.*, during each evolutionary step, only the best candidature are kept.

Firstly, we utilize existing individual adversarial objectives (*i.e.*, conventional GANs) to perform the adversarial training process. We train each method 50K iterations and report the KDE plots in Fig. 3. Meanwhile, the average and the best MMD values which running each method 10 times are reported in Table II. The results show that all of the individual adversarial objectives suffer from mode collapse to a greater or lesser degree. Then, we set hyper-parameter γ as 0 and test the proposed E-GAN, which means the diversity fitness score was not considered during the training process. The results show that the evolutionary framework still troubles with the mode collapse. However, when the diversity fitness score

³We obtain both 2D distributions and network architectures from the code provided in [44].

TABLE III

INCEPTION SCORES AND FIDS WITH UNSUPERVISED IMAGE GENERATION ON CIFAR-10. THE METHOD WITH HIGHER INCEPTION SCORE OR LOWER FID IMPLIES THE GENERATED DISTRIBUTION IS CLOSER TO THE TARGET ONE. [†] [22], [‡] [6].

Methods	Inception score	FID
Real data	11.24 \pm .12	7.8
-Standard CNN-		
(ours) E-GAN-GP ($\mu = 1$)	7.13 \pm .07	33.2
(ours) E-GAN-GP ($\mu = 2$)	7.23 \pm .08	31.6
(ours) E-GAN-GP ($\mu = 4$)	7.32 \pm .09	29.8
(ours) E-GAN-GP ($\mu = 8$)	7.34 \pm .07	27.3
(ours) E-GAN ($\mu = 1$, without GP)	6.98 \pm .09	36.2
DCGAN (without GP) [†]	6.64 \pm .14	-
GAN-GP [‡]	6.93 \pm .08	37.7
WGAN-GP [‡]	6.68 \pm .06	40.2

is considered in the selection stage (in this experiment, we set $\gamma = 0.5$), the mode collapse issue is largely suppressed and the trained generator can more accurately fit the target distributions. This demonstrates, our diversity fitness score has the capability of measuring sample diversity of updated generators and further suppress the mode collapse problem.

D. CIFAR-10 and Training Stability

In the proposed E-GAN, we utilize the evolutionary algorithm with different mutations (*i.e.*, different updating strategies) to optimize generator(s) $\{G\}$. To demonstrate the advantages of the proposed evolutionary algorithm over existing two-player adversarial training strategies (*i.e.*, updating generator with a single objective), we train these methods on CIFAR-10 and plot inception scores [43] over the training process. For a fair comparison, we did not introduce the gradient penalty norm into our E-GAN and set the parents




Fig. 4. Experiments on the CIFAR-10 dataset. CIFAR-10 inception score over generator iterations (left), over wall-clock time (middle), and the graph of selected mutations in the E-GAN training process (right).

TABLE IV
FIDS WITH UNSUPERVISED IMAGE GENERATION ON LSUN BEDROOM DATASET. THE METHOD WITH LOWER FID IMPLIES THE GENERATED DISTRIBUTION IS CLOSER TO THE TARGET ONE.

Methods	FID			
	-Baseline-	-Weak G-	-Weak D-	-Weak both-
DCGAN	43.7	187.3	410.6	82.7
LSGAN	46.3	452.9	423.1	126.2
WGAN	51.1	113.6	129.2	115.7
WGAN-GP	38.5	66.7	385.8	73.2
(ours) E-GAN ($\mu = 1$, without GP)	34.2	63.3	64.8	71.9
(ours) E-GAN ($\mu = 4$, without GP)	29.7	59.1	55.2	60.9

number μ as 1. Moreover, the same network architecture is used for all methods.

As shown in Fig. 4-left, E-GAN can get higher inception score with less training steps. Meanwhile, E-GAN also shows comparable stability when it goes to convergence. By comparison, conventional GAN objectives expose their different limitations, such as instability at convergence (GAN-Heuristic), slow convergence (GAN-Least square) and invalid (GAN-minimax). In addition, we employ the proposed fitness function (*i.e.*, Eq. (9)) as generator's objective function, and find its performance is also inferior to E-GAN. This experiment further demonstrates the advantages of the proposed evolutionary framework. Through creating and selecting from multiple candidates, the evolutionary framework can leverage strengths of different objective functions (*i.e.*, different distances) to accelerate the training process and improve the generative performance. Based on the evolutionary framework, E-GAN not only overcomes the inherent limitations of these individual adversarial objectives, but it also outperforms other GANs (the WGAN and its improved variant WGAN-GP). Furthermore, when we only keep one parent during each evolutionary step, E-GAN achieves comparable convergence speed in terms of wall-clock time (Fig. 4-middle). During training E-GAN, we recorded the selected objective in each step (Fig. 4-right). At the beginning of training, the heuristic objective and the least-square objective are selected more frequently than the minimax objective. It may due to the fact that the minimax objective

is hard to provide effective gradients (*i.e.*, vanishing gradient) when the discriminator can easily recognize generated samples. Along with the generator approaching convergence (after 20K steps), ever more minimax objectives are employed, yet the number of selected heuristic objectives is falling. As aforementioned, the minus JSDs of the heuristic objective may tend to push the generated distribution away from target data distribution and lead to training instability. However, in E-GAN, beyond the heuristic objective, we have other options of objective, which improves the stability at convergence.

Furthermore, we discussed the relationship between the survived parents' number and generative performance. As shown in the Table III, both the Inception score and FID are utilized to evaluate the generative performance of learned generators. Firstly, compared with the basic E-GAN (*i.e.*, E-GAN, $\mu = 1$, without GP), adding the gradient penalty (GP) norm to optimize the discriminator indeed improves generative performance. Then, we preserved multiple parents during the E-GAN training process and measured their scores at the end of training. We can easily observe that the generative performance becomes better accompanying with keeping more parents during the training. This further demonstrates the proposed evolutionary learning paradigm could suppress the unstable and large-scale optimization problems of GANs.

Theoretically, if we regard updating and evaluating a child G as an operation and define mutations number as n , keeping p parents will cost $O(np)$ operations in each iteration. Com-

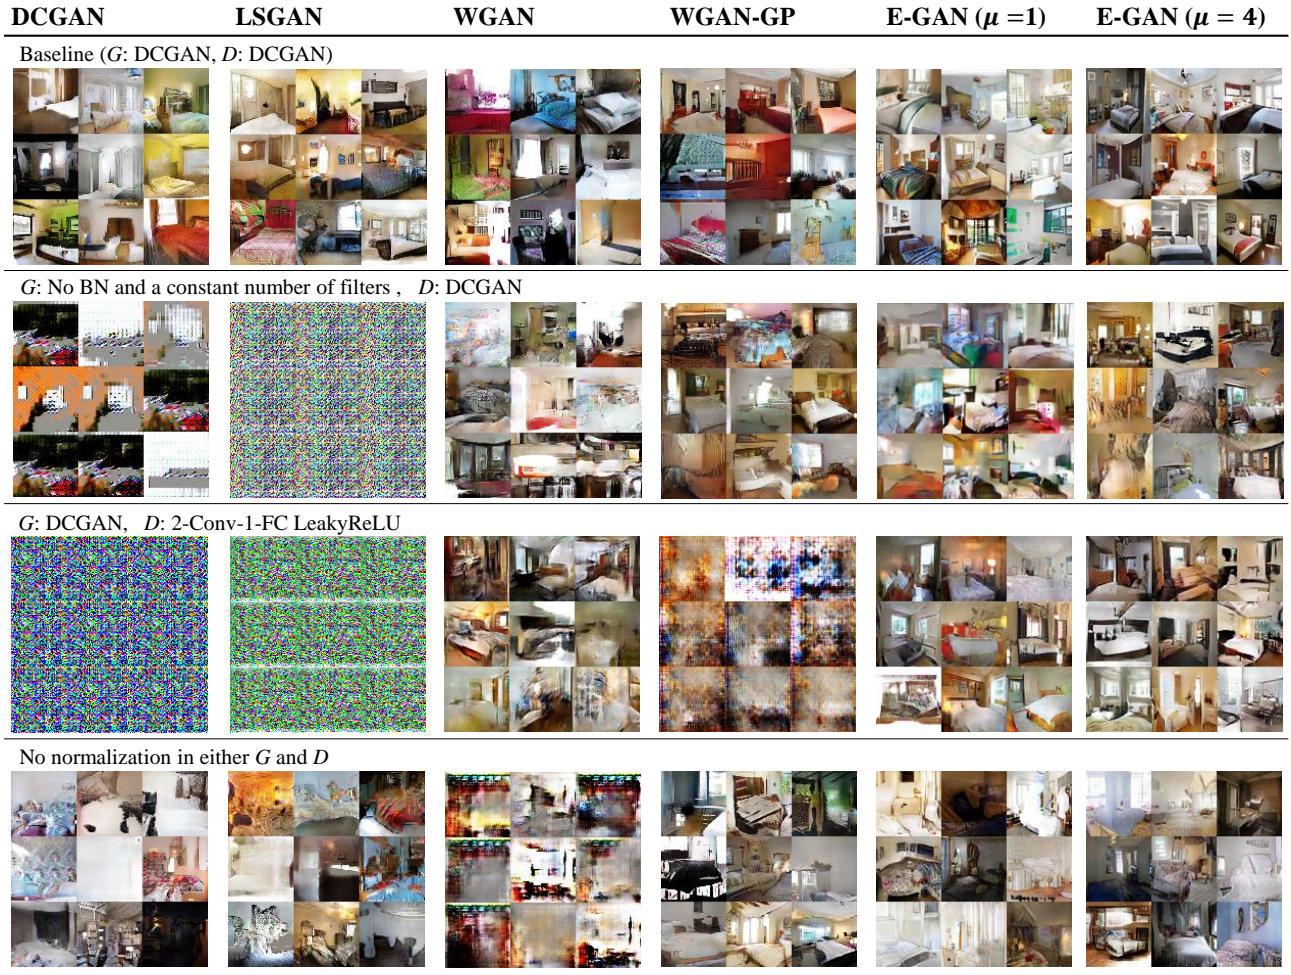


Fig. 5. Experiments of architecture robustness. Different GAN architectures, which correspond to different training challenges, trained with six different GAN methods (or settings). The proposed E-GAN achieved promising performance under all architecture settings.

paring with traditional GANs, our evolutionary framework would cost more time in each iteration. However, since the evolutionary strategy always preserves the well-performing off-spring, to achieve the same generative performance, E-GAN usually spends less training steps (Fig. 4-left). Overall, keeping one parent during each evolutionary step will only slightly reduce the time-efficiency but with better performance (Fig. 4-middle). Yet, accompanied by increasing p , although the generative performance can be further improved, it will also cost more time on training the E-GAN model. Here, if we regard the parents number p as a hyper-parameter of our algorithm, we found setting its value less than or equal to 4 is a preferable choice. Within this interval, we can easily improve the generative performance by sacrificing affordable computation cost. If we continually increase the number of survivors, the generative performance can only be improved mildly yet largely reduce the training efficiency. In practice, we need to further balance the algorithms efficiency and performance according to different situations.

E. LSUN and Architecture Robustness

The architecture robustness is another advantage of E-GAN. To demonstrate the training stability of our method, we



Fig. 6. Generated bedroom images on 128×128 LSUN bedrooms.

train different network architectures on the LSUN bedroom dataset [68] and compare with several existing works. In addition to the baseline DCGAN architecture, we choose three additional architectures corresponding to different train-

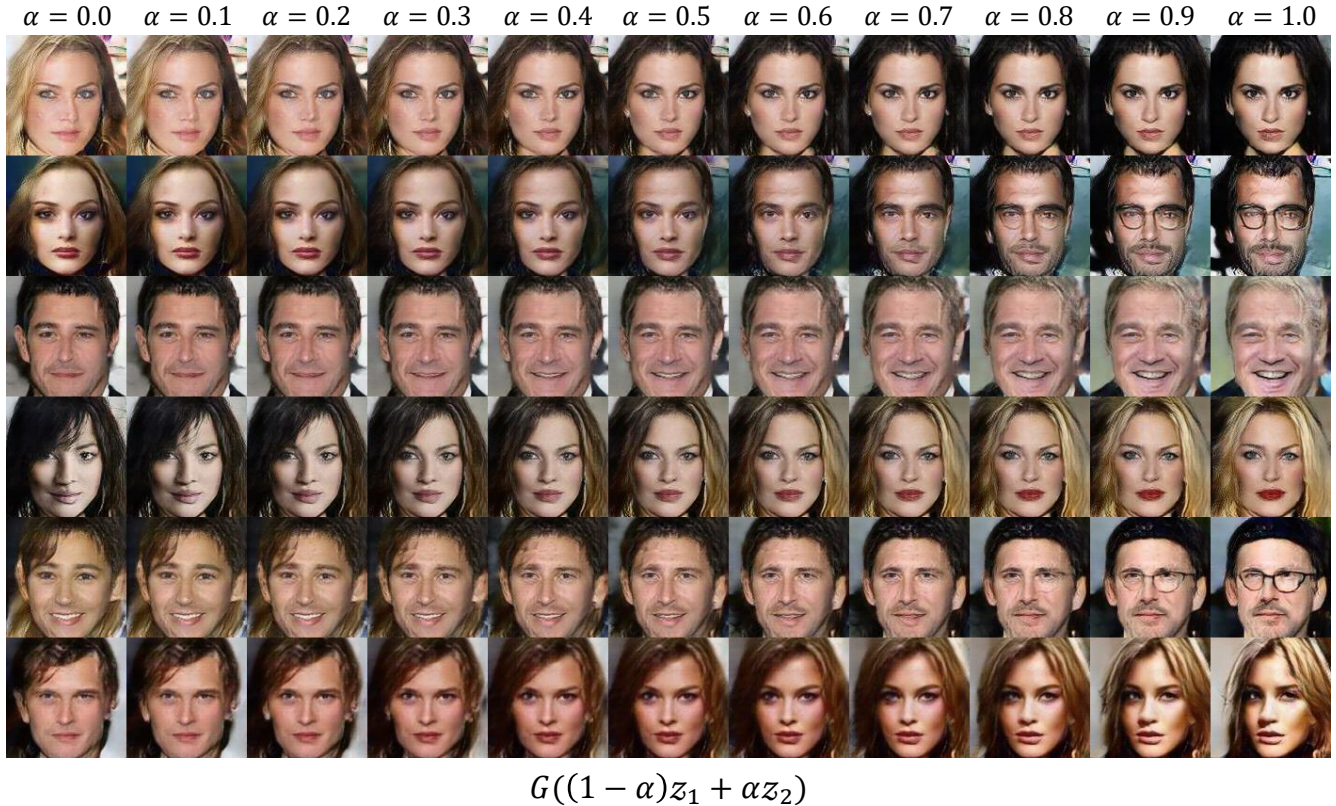


Fig. 7. Interpolating in latent space. For selected pairs of generated images from a well-trained E-GAN model, we record their latent vectors z_1 and z_2 . Then, samples between them are generated by linear interpolation between these two vectors.

ing challenges: (1) limiting the recognition capability of the discriminator D , *i.e.*, 2-Conv-1-FC LeakyReLU discriminator (abbreviated as weak D); (2) limiting the expression capability of the generator G , *i.e.*, no batchnorm and a constant number of filters in the generator (weak G); (3) reducing the network capability of the generator and discriminator together, *i.e.*, remove the BN in both the generator G and discriminator D (weak both). For each architecture, we test six different methods (or settings): DCGAN, LSGAN, original WGAN (with weight clipping), WGAN-GP (with gradient penalty), our E-GAN ($\mu = 1$), and E-GAN ($\mu = 4$). For each method, we used the default configurations recommended in the respective studies (these methods are summarized in [44]) and train each model for 100K iterations. Some generated samples are reported in Fig. 5, and the quantitative results (*i.e.*, FID) are listed in Table IV. Through observation, we find that all of these GAN methods achieved promising performance with the baseline architecture. For DCGAN and LSGAN, when the balance between the generator and discriminator is broken (*i.e.*, only one of them is limited), these two methods are hard to generate any reasonable samples. Meanwhile, we find the performance of the standard WGAN (with weight clipping) is mostly decided by the generator G . When we limit G 's capability, the generative performance largely reduced. As regarded the WGAN-GP, we find that the generative performance may mainly depends on the discriminator (or critic). Our E-GAN achieved promising results under all architecture settings. Moreover, we again demonstrated that the model performance is growing with the number of survived parents.



Fig. 8. Generated human face images on 128×128 CelebA dataset.

Furthermore, we trained E-GAN to generate higher resolution (128×128) bedroom images (Fig. 6). Observing generated images, we demonstrate that E-GAN can be trained to generate diversity and high-quality images from the target data distribution.

F. CelebA and Space Continuity

Besides the LSUN bedroom dataset, we also train our E-GAN using the aligned faces from CelebA dataset. Since humans excel at identifying facial flaws, generating high-quality human face images is challenging. Similar to generating bedrooms, we employ the same architectures to generate 128×128 RGB human face images (Fig. 8). In addition,

given a well-trained generator, we evaluate the performance of the embedding in the latent space of noisy vectors z . In Fig. 7, we first select pairs of generated faces and record their corresponding latent vectors z_1 and z_2 . The two images in one pair have different attributes, such as gender, expression, hairstyle, and age. Then, we generate novel samples by linear interpolating between these pairs (*i.e.*, corresponding noisy vectors). We find that these generated samples can seamlessly change between these semantically meaningful face attributes. This experiment demonstrates that generator training does not merely memorize training samples but learns a meaningful projection from latent noisy space to face images. Meanwhile, it also shows that the generator trained by E-GAN does not suffer from mode collapse, and shows great space continuity. Overall, during the GAN training process, the training stability is easily influenced by ‘bad’ updating, which could lead the generated samples to low quality or lacking diversity, while the proposed evolutionary mechanism largely avoids undesired updating and promote the training to an ideal direction.

V. CONCLUSION

In this paper, we present an evolutionary GAN framework (E-GAN) for training deep generative models. To reduce training difficulties and improve generative performance, we devise an evolutionary algorithm to evolve a population of generators to adapt to the dynamic environment (*i.e.*, the discriminator D). In contrast to conventional GANs, the evolutionary paradigm allows the proposed E-GAN to overcome the limitations of individual adversarial objectives and preserve the well-performing offspring after each iteration. Experiments show that E-GAN improves the training stability of GAN models and achieves convincing performance in several image generation tasks. In this work, we mainly contribute to improving the image generation performance. More generation tasks will be considered in future works, such as video generation [33] and text generation [77].

REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems (NIPS)*, 2014, pp. 2672–2680.
- [2] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, “Infogan: Interpretable representation learning by information maximizing generative adversarial nets,” in *Advances in Neural Information Processing Systems (NIPS)*, 2016, pp. 2172–2180.
- [3] Z. Gan, L. Chen, W. Wang, Y. Pu, Y. Zhang, H. Liu, C. Li, and L. Carin, “Triangle generative adversarial networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [4] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas, “Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks,” in *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [5] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of gans for improved quality, stability, and variation,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [6] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, “Spectral normalization for generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [7] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] C. Wang, C. Wang, C. Xu, and D. Tao, “Tag disentangled generative adversarial networks for object image re-rendering,” in *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017, pp. 2901–2907.
- [9] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [10] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, “High-resolution image synthesis and semantic manipulation with conditional gans,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] C. Vondrick, H. Pirsiavash, and A. Torralba, “Generating videos with scene dynamics,” in *Advances in Neural Information Processing Systems (NIPS)*, 2016, pp. 613–621.
- [12] C. Finn, I. Goodfellow, and S. Levine, “Unsupervised learning for physical interaction through video prediction,” in *Advances in Neural Information Processing Systems (NIPS)*, 2016, pp. 64–72.
- [13] C. Vondrick and A. Torralba, “Generating the future with adversarial transformers,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [14] Y. Zhang, Z. Gan, and L. Carin, “Generating text via adversarial training,” in *NIPS workshop on Adversarial Training*, 2016.
- [15] W. Fedus, I. Goodfellow, and A. M. Dai, “Maskgan: Better text generation via filling in the _,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [16] J. Lu, A. Kannan, J. Yang, D. Parikh, and D. Batra, “Best of both worlds: Transferring knowledge from discriminative learning to a generative visual dialog model,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [17] X. Lan, A. J. Ma, and P. C. Yuen, “Multi-cue visual tracking using robust feature-level fusion based on joint sparse representation,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1194–1201.
- [18] X. Chen, C. Xu, X. Yang, L. Song, and D. Tao, “Gated-gan: Adversarial gated networks for multi-collection style transfer,” *IEEE Transactions on Image Processing*, vol. 28, no. 2, pp. 546–560, 2019.
- [19] S. Arora, R. Ge, Y. Liang, T. Ma, and Y. Zhang, “Generalization and equilibrium in generative adversarial nets (GANs),” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 224–232.
- [20] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, “Least squares generative adversarial networks,” in *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [21] J. Zhao, M. Mathieu, and Y. LeCun, “Energy-based generative adversarial network,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [22] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- [23] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 214–223.
- [24] M. Arjovsky and L. Bottou, “Towards principled methods for training generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [25] V. Nagarajan and J. Z. Kolter, “Gradient descent gan optimization is locally stable,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [26] C. Qian, J.-C. Shi, K. Tang, and Z.-H. Zhou, “Constrained monotone k-submodular function maximization using multi-objective evolutionary algorithms with theoretical guarantee,” *IEEE Transactions on Evolutionary Computation*, 2017.
- [27] S. He, G. Jia, Z. Zhu, D. A. Tennant, Q. Huang, K. Tang, J. Liu, M. Musolesi, J. K. Heath, and X. Yao, “Cooperative co-evolutionary module identification with application to cancer disease module discovery,” *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 6, pp. 874–891, 2016.
- [28] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, “The expressive power of parameterized quantum circuits,” *arXiv preprint arXiv:1810.11922*, 2018.
- [29] L. M. Antonio and C. A. C. Coello, “Coevolutionary multi-objective evolutionary algorithms: A survey of the state-of-the-art,” *IEEE Transactions on Evolutionary Computation*, 2017.
- [30] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, Q. Le, and A. Kurakin, “Large-scale evolution of image classifiers,” in *Proceedings*

- of the 34th International Conference on Machine Learning (ICML), 2017, pp. 214–223.
- [31] T. Salimans, J. Ho, X. Chen, and I. Sutskever, “Evolution strategies as a scalable alternative to reinforcement learning,” *arXiv preprint arXiv:1703.03864*, 2017.
 - [32] I. Goodfellow, “Nips 2016 tutorial: Generative adversarial networks,” *arXiv preprint arXiv:1701.00160*, 2016.
 - [33] S. Tulyakov, M.-Y. Liu, X. Yang, and J. Kautz, “Mocogan: Decomposing motion and content for video generation,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [34] Y. Song, C. Ma, X. Wu, L. Gong, L. Bao, W. Zuo, C. Shen, R. W. Lau, and M.-H. Yang, “Vital: Visual tracking via adversarial learning,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [35] X. Lan, S. Zhang, P. C. Yuen, and R. Chellappa, “Learning common and feature-specific patterns: a novel multiple-sparse-representation-based tracker,” *IEEE Transactions on Image Processing*, vol. 27, no. 4, pp. 2022–2037, 2018.
 - [36] X. Lan, A. J. Ma, P. C. Yuen, and R. Chellappa, “Joint sparse representation and robust feature-level fusion for multi-cue visual tracking,” *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5826–5841, 2015.
 - [37] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, “Adversarial discriminative domain adaptation,” in *Computer Vision and Pattern Recognition (CVPR)*, vol. 1, no. 2, 2017, p. 4.
 - [38] Z. Qiu, Y. Pan, T. Yao, and T. Mei, “Deep semantic hashing with generative adversarial networks,” in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2017, pp. 225–234.
 - [39] E. Yang, T. Liu, C. Deng, and D. Tao, “Adversarial examples for hamming space search,” *IEEE transactions on cybernetics*, 2018.
 - [40] E. Yang, C. Deng, W. Liu, X. Liu, D. Tao, and X. Gao, “Pairwise relationship guided deep hashing for cross-modal retrieval,” in *AAAI*, 2017, pp. 1618–1625.
 - [41] J. Donahue, P. Krähenbühl, and T. Darrell, “Adversarial feature learning,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
 - [42] V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. Courville, “Adversarially learned inference,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
 - [43] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, “Improved techniques for training gans,” in *Advances in Neural Information Processing Systems (NIPS)*, 2016, pp. 2234–2242.
 - [44] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved training of wasserstein gans,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
 - [45] T. D. Nguyen, T. Le, H. Vu, and D. Phung, “Dual discriminator generative adversarial nets,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
 - [46] I. Durugkar, I. Gemp, and S. Mahadevan, “Generative multi-adversarial networks,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
 - [47] B. Neyshabur, S. Bhojanapalli, and A. Chakrabarti, “Stabilizing gan training with multiple random projections,” *arXiv preprint arXiv:1705.07831*, 2017.
 - [48] I. O. Tolstikhin, S. Gelly, O. Bousquet, C.-J. Simon-Gabriel, and B. Schölkopf, “Adagan: Boosting generative models,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 5424–5433.
 - [49] A. Ghosh, V. Kulharia, V. P. Namboodiri, P. H. Torr, and P. K. Dokania, “Multi-agent diverse generative adversarial networks,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
 - [50] Q. Hoang, T. D. Nguyen, T. Le, and D. Phung, “Mgan: Training generative adversarial nets with multiple generators,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
 - [51] K. A. De Jong, *Evolutionary computation: a unified approach*. MIT press, 2006.
 - [52] H.-L. Liu, L. Chen, Q. Zhang, and K. Deb, “Adaptively allocating search effort in challenging many-objective optimization problems,” *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 3, pp. 433–448, 2018.
 - [53] Y. Wang, M. Zhou, X. Song, M. Gu, and J. Sun, “Constructing cost-aware functional test-suites using nested differential evolution algorithm,” *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 3, pp. 334–346, 2018.
 - [54] P. Yang, K. Tang, and X. Yao, “Turning high-dimensional optimization into computationally expensive optimization,” *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 1, pp. 143–156, 2018.
 - [55] D.-C. Dang, T. Friedrich, T. Kötzing, M. S. Krejca, P. K. Lehre, P. S. Oliveto, D. Sudholt, and A. M. Sutton, “Escaping local optima using crossover with emergent diversity,” *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 3, pp. 484–497, 2018.
 - [56] A. E. Eiben and J. Smith, “From evolutionary computation to the evolution of things,” *Nature*, vol. 521, no. 7553, p. 476, 2015.
 - [57] R. Miikkulainen, J. Liang, E. Meyerson, A. Rawal, D. Fink, O. Francon, B. Raju, A. Navruzyan, N. Duffy, and B. Hodjat, “Evolving deep neural networks,” *arXiv preprint arXiv:1703.00548*, 2017.
 - [58] S. R. Young, D. C. Rose, T. P. Karnowski, S.-H. Lim, and R. M. Patton, “Optimizing deep learning hyper-parameters through an evolutionary algorithm,” in *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments*. ACM, 2015, p. 4.
 - [59] X. Yao, “Evolving artificial neural networks,” *Proceedings of the IEEE*, vol. 87, no. 9, pp. 1423–1447, 1999.
 - [60] S. Lander and Y. Shang, “Evoac—a new evolutionary method for training autoencoders for deep learning networks,” in *Computer Software and Applications Conference (COMPSAC), 2015 IEEE 39th Annual*, vol. 2. IEEE, 2015, pp. 790–795.
 - [61] Y. Wang, C. Xu, J. Qiu, C. Xu, and D. Tao, “Towards evolutionary compression,” *arXiv preprint arXiv:1707.08005*, 2017.
 - [62] R. S. Olson, N. Bartley, R. J. Urbanowicz, and J. H. Moore, “Evaluation of a tree-based pipeline optimization tool for automating data science,” in *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. ACM, 2016, pp. 485–492.
 - [63] A. Rosales-Pérez, S. García, J. A. Gonzalez, C. A. C. Coello, and F. Herrera, “An evolutionary multiobjective model and instance selection for support vector machines with pareto-based ensembles,” *IEEE Transactions on Evolutionary Computation*, vol. 21, no. 6, pp. 863–877, 2017.
 - [64] M. Mirza and S. Osindero, “Conditional generative adversarial nets,” *arXiv preprint arXiv:1411.1784*, 2014.
 - [65] J. Liang, J. Yang, H.-Y. Lee, K. Wang, and M.-H. Yang, “Sub-gan: An unsupervised generative model via subspaces,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 698–714.
 - [66] O. Kramer, *Machine learning for evolution strategies*. Springer, 2016, vol. 20.
 - [67] A. Krizhevsky, “Learning multiple layers of features from tiny images,” 2009.
 - [68] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao, “Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop,” *arXiv preprint arXiv:1506.03365*, 2015.
 - [69] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *The IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 3730–3738.
 - [70] W. Fedus, M. Rosca, B. Lakshminarayanan, A. M. Dai, S. Mohamed, and I. Goodfellow, “Many paths to equilibrium: Gans do not need to decrease adivergence at every step,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
 - [71] A. Smola, A. Gretton, L. Song, and B. Schölkopf, “A hilbert space embedding for distributions,” in *International Conference on Algorithmic Learning Theory*. Springer, 2007, pp. 13–31.
 - [72] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, “A kernel two-sample test,” *Journal of Machine Learning Research*, vol. 13, no. Mar, pp. 723–773, 2012.
 - [73] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a local nash equilibrium,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 6629–6640.
 - [74] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
 - [75] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, “Self-attention generative adversarial networks,” *arXiv preprint arXiv:1805.08318*, 2018.
 - [76] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein, “Unrolled generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
 - [77] N. Duan, D. Tang, P. Chen, and M. Zhou, “Question generation for question answering,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017, pp. 866–874.



Chaoyue Wang is research associate in Machine Learning and Computer Vision at the School of Computer Science, The University of Sydney. He received a bachelor degree from Tianjin University (TJU), China, and a Ph.D. degree from University of Technology Sydney (UTS), Australia. His research interests mainly include machine learning, deep learning, and generative models. He received the Distinguished Student Paper Award in the 2017 International Joint Conference on Artificial Intelligence (IJCAI).



Chang Xu is Lecturer in Machine Learning and Computer Vision at the School of Computer Science, The University of Sydney. He obtained a Bachelor of Engineering from Tianjin University, China, and a Ph.D. degree from Peking University, China. While pursuing his PhD degree, Chang received fellowships from IBM and Baidu. His research interests lie in machine learning, data mining algorithms and related applications in artificial intelligence and computer vision, including multi-view learning, multi-label learning, visual search and face recognition.

His research outcomes have been widely published in prestigious journals and top-tier conferences.



Xin Yao (F'03) obtained his PhD in 1990 from the University of Science and Technology of China (USTC), MSc in 1985 from North China Institute of Computing Technologies and BSc in 1982 from USTC. He is a Chair Professor of Computer Science at the Southern University of Science and Technology, Shenzhen, China, and a part-time Professor of Computer Science at the University of Birmingham, UK. He is an IEEE Fellow and was a Distinguished Lecturer of IEEE Computational Intelligence Society (CIS). His major research interests include evolutionary computation, ensemble learning, and their applications to software engineering.

His paper on evolving artificial neural networks won the 2001 IEEE Donald G. Fink Prize Paper Award. He also won 2010, 2016 and 2017 IEEE Transactions on Evolutionary Computation Outstanding Paper Awards, 2011 IEEE Transactions on Neural Networks Outstanding Paper Award, and many other best paper awards. He received the prestigious Royal Society Wolfson Research Merit Award in 2012 and the IEEE CIS Evolutionary Computation Pioneer Award in 2013. He was the President (2014-15) of IEEE CIS and the Editor-in-Chief (2003-08) of IEEE Transactions on Evolutionary Computation.



Dacheng Tao (F'15) is Professor of Computer Science and ARC Laureate Fellow in the School of Information Technologies and the Faculty of Engineering and Information Technologies, and the Inaugural Director of the UBTECH Sydney Artificial Intelligence Centre, at the University of Sydney. He mainly applies statistics and mathematics to Artificial Intelligence and Data Science. His research results have expounded in one monograph and 200+ publications at prestigious journals and prominent conferences, such as IEEE T-PAMI, T-IP, T-NNLS, IJCV, JMLR, NIPS, ICML, CVPR, ICCV, ECCV, ICDM; and ACM SIGKDD, with several best paper awards, such as the best theory/algorithm paper runner up award in IEEE ICDM07, the best student paper award in IEEE ICDM13, the distinguished paper award in the 2018 IJCAI, the 2014 ICDM 10-year highest-impact paper award, and the 2017 IEEE Signal Processing Society Best Paper Award. He received the 2015 Austrian Scopus-Eureka Prize and the 2018 IEEE ICDM Research Contributions Award. He is a Fellow of the Australian Academy of Science, AAAS, IEEE, IAPR, OSA and SPIE.

His research results have expounded in one monograph and 200+ publications at prestigious journals and prominent conferences, such as IEEE T-PAMI, T-IP, T-NNLS, IJCV, JMLR, NIPS, ICML, CVPR, ICCV, ECCV, ICDM; and ACM SIGKDD, with several best paper awards, such as the best theory/algorithm paper runner up award in IEEE ICDM07, the best student paper award in IEEE ICDM13, the distinguished paper award in the 2018 IJCAI, the 2014 ICDM 10-year highest-impact paper award, and the 2017 IEEE Signal Processing Society Best Paper Award. He received the 2015 Austrian Scopus-Eureka Prize and the 2018 IEEE ICDM Research Contributions Award. He is a Fellow of the Australian Academy of Science, AAAS, IEEE, IAPR, OSA and SPIE.