Generative Adversarial Networks: Introduction and Outlook

Kunfeng Wang, *Member, IEEE*, Chao Gou, Yanjie Duan, Yilun Lin, Xinhu Zheng, and Fei-Yue Wang, *Fellow, IEEE*

***Abstract*—Recently, generative adversarial networks (GANs) have become a research focus of artificial intelligence. Inspired by two-player zero-sum game, GANs comprise a generator and a discriminator, both trained under the adversarial learning idea. The goal of GANs is to estimate the potential distribution of real data samples and generate new samples from that distribution. Since their initiation, GANs have been widely studied due to their enormous prospect for applications, including image and vision computing, speech and language processing, etc. In this review paper, we summarize the state of the art of GANs and look into the future. Firstly, we survey GANs' proposal background, theoretic and implementation models, and application fields. Then, we discuss GANs' advantages and disadvantages, and their development trends. In particular, we investigate the relation between GANs and parallel intelligence, with the conclusion that GANs have a great potential in parallel systems research in terms of virtual-real interaction and integration. Clearly, GANs can provide substantial algorithmic support for parallel intelligence.**

***Index Terms*—ACP approach, adversarial learning, generative adversarial networks (GANs), generative models, parallel intelligence, zero-sum game.**

I. Introduction

GENERATIVE adversarial networks (GANs) are a powerful class of generative models introduced in 2014 by Goodfellow *et al.* [1]. The basic principle of GANs is inspired by two-player zero-sum game, in which the total gains of two players are zero, and each player's gain or loss of utility is exactly balanced by the loss or gain of the utility of another player. GANs often comprise a generator and a discriminator that learn simultaneously. The generator tries to capture the potential distribution of real samples, and generates new data samples. The discriminator is often a binary classifier, discriminating real samples from the generated samples as accurately as possible. Both the generator and the discriminator can adopt the structure of currently popular deep neural networks [2], [3]. The optimization process of GANs is a minimax game process, and the optimization goal is to reach Nash equilibrium [4], where the generator is considered to have captured the distribution of real samples.

Under the boom of artificial intelligence, the proposal of GANs satisfies the research and application requirements of many fields, and injects fresh impetus to the development of the related fields. GANs have become a hot research topic in artificial intelligence. Yann LeCun, in a recent lecture on unsupervised learning, calls adversarial networks "the coolest idea in machine learning in the last twenty years". Nowadays, the image and vision field receives the most attention of GANs researchers. It is now possible using GANs to generate photorealistic object images such as birds and faces, generate indoor or outdoor scenes, translate images from a source domain to the target domain, generate high-definition images from low-definition images, and so on [5]. Besides, GANs have been introduced into the study of other artificial intelligence subfields, including speech and language processing [6], [7], malware detection [8], and chess game program [9].

This paper surveys the state of the art of GANs and looks into their future. Section II introduces GANs' proposal background. Section III describes GANs' theoretic and implementation models, including GANs' basic principle, learning method, and GAN variants. Section IV presents some typical applications of GANs in artificial intelligence. Section V discusses GANs' advantages and disadvantages, and their development trends. In particular, the relation between GANs and parallel intelligence is investigated in Section V. Finally, the concluding remarks are made in Section VI.

II. Proposal Background

We introduce the proposal background of GANs in this section, in order to make readers have a better understanding of GANs' research progress and application fields.

Manuscript received April 9, 2017; accepted May 17, 2017. This work was supported by the National Natural Science Foundation of China (61533019, 71232006, 91520301). Recommended by Associate Editor Derong Liu. *(Corresponding author: Fei-Yue Wang.)*

Citation: K. F. Wang, C. Gou, Y. J. Duan, Y. L. Lin, X. H. Zheng, and F.-Y. Wang, "Generative adversarial networks: introduction and outlook," *IEEE/CAA J. of Autom. Sinica*, vol. 4, no. 4, pp. 588−598, Oct. 2017.

K. F. Wang is with the State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China, and also with Qingdao Academy of Intelligent Industries, Qingdao 266000, China (e-mail: kunfeng.wang@ia.ac.cn).

C. Gou, Y. J. Duan, and Y. L. Lin are with the State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China, and also with University of Chinese Academy of Sciences, Beijing 100049, China (e-mail: gouchao2012@ia.ac.cn; duanyanjie2012@ia.ac.cn; linyilun2014@ia.ac.cn).

X. H. Zheng is with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis MN 55414, USA (e-mail: zheng473@umn.edu).

F. -Y. Wang is with the State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China, and also with the Research Center for Computational Experiments and Parallel Systems Technology, National University of Defense Technology, Changsha 410073, China (e-mail: feiyue.wang@ia.ac.cn).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JAS.2017.7510583

A. Boom of Artificial Intelligence

In recent years, with the increase in computing power and the emergence of big data in various industries, artificial intelligence (AI) has gained rapid development and wide application. Both the researchers' attention to AI and the public's desire for AI utility are improving unprecedentedly [2], [10]. It is generally believed that AI can be divided into two stages: perception and cognition. In the perception stage, AI systems receive physical signals (such as video and audio signals) from the real world and make discriminations about the signals. Related research areas include image recognition, speech recognition, and so on. In the cognition stage, AI systems should have a certain understanding of the nature of the world rather than make only discriminations mechanically. Based on our research experience, we think that AI has four levels including discrimination, generation, understanding, and creation and application, as shown in Fig. 1. On one hand, these levels are interrelated and reinforce each other. On the other hand, there are big gaps between different levels, awaiting new research breakthroughs.

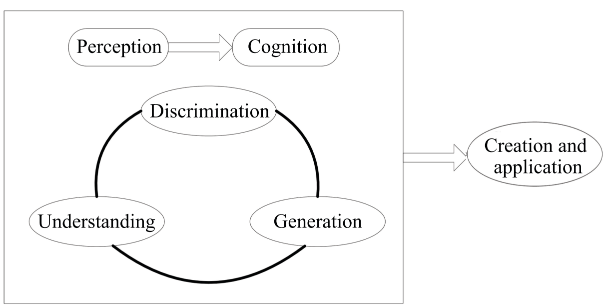

Fig. 1. The research levels of Artificial Intelligence.

Both the widely believed two stages of AI and our summarized four levels of AI involve a common topic: understanding. Nevertheless, understanding cannot be measured directly because it is an internal characteristic both for humans and for AI. Understanding can only be inferred indirectly from other aspects. Although how to measure the understanding level of AI remains unresolved, a famous scholar Richard Feynman has a saying "What I cannot create, I do not understand", indicating that to some extent the ability of machines to make things reflects the level of machines to understand things. GAN as a typical generative model can generate data samples with its generator. This ability reflects its understanding of things up to a certain degree. Thus, it is expected that GANs can deepen the AI research.

B. Accumulation of Generative Models

Generative models play an essential role in the AI field, even the generation methods alone have great research significance. Generative methods and discriminative methods are two branches of supervised learning in machine learning. The generative models are the models obtained by learning with the generative methods. The generative methods involve distribution hypothesis and parameter estimation, and can sample new data from the estimated models. We argue that generative models have two research perspectives: humans understand data and machines understand data.

From the perspective of "humans understand data", a typical approach involves assuming the distribution of explicit or latent variables, and then using real data to fit the distribution parameters or train the model containing the distribution. After that, a new sample is generated using the learned distribution or model. The methods belonging to this class of generative models include maximum likelihood estimation, approximate inference [11], [12], and Markov chain method [13]−[15]. The model learned from this perspective has a distribution that humans can understand, but has limitations for learning machines. For example, the maximum likelihood estimation is conducted on real data samples, and the parameters are updated directly according to the data samples, leading to an overly smooth generative model. The generative model obtained by approximate inference can only approach the lower bound of the objective function instead of directly approaching the objective function, due to the difficulty in solving the objective function. The Markov chain method can be used for training generative models and generating new samples, but its computational complexity is extremely high.

From the perspective of "machines understand data", the generative model does not directly estimate or fit the data distribution. Instead, it generates data samples from the distribution without explicit hypothesis [16], and use the generated samples to modify the model. The resulting generative model is less interpretable to humans, but the newly generated samples are understandable to humans. It is conjectured that machines understand data in a way that human cannot understand explicitly, but generate new data that human can understand. Prior to GANs, the generative models built from the perspective of "machines understand data" generally need to be trained using Markov chain, which has low efficiency and limits their systematic applications.

Before the proposal of GANs, the generative models already have certain research accumulations. However, the limitations existing in model training and data generation are really barriers of generative models. To realize the four levels of AI, it is necessary to design a new paradigm of generative models to break through the existing barriers.

C. Deepening of Neural Networks

In the past decade, with the great success of deep learning [17], [18] in various fields, the research on neural networks revives again. Due to the increase in computing power and data scale, neural networks have been overcoming the difficulty in parameter training, and are widely used to solve complex nonlinear problems. For example, deep learning has achieved a breakthrough effect in image classification [19], [20], and significantly improved the accuracy of speech recognition [21]. It has also been successfully applied in natural language processing and understanding [22]. The success of neural networks is closely related to the model characteristics. In terms of model training, neural networks can use the general backpropagation algorithm, and the training process is easy to realize. In terms of model structure, the structural design of

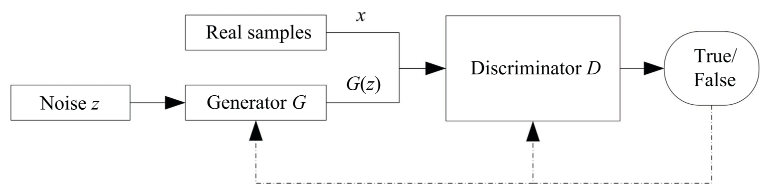

Fig. 2. Computation procedure and structure of GAN.

neural networks is free and flexible with few restrictions. In terms of model capacity, neural networks can approximate any function theoretically, and have a wide range of applications. With the rapid development of computing devices, neural networks with more parameters can be trained with higher speed, further promoting the popularity of neural networks.

D. Success of Adversarial Idea

From machine learning to AI, the adversarial idea has been successfully introduced. Both game and competition contain the adversarial idea. Game-theoretic machine learning [23] combines game theory with machine learning, models humans' dynamic strategy by game theory and optimizes advertisement auction mechanism, then proves the effectiveness of the approach through experiments. The more recent event that AlphaGo [24] defeats human masters triggers public interest in AI. The intermediate version of AlphaGo uses two networks fighting with each other in the process of training policy network, obtains the game state, policy and the corresponding return, and takes the expectation function containing the game return as the maximizing objective. In the study of neural network, researchers have used two neural networks to compete with each other in the training process [25], encouraging the hidden nodes in the network to be statistically independent, which is used as a regularization factor in the training process. There are researchers [26], [27] using the adversarial idea to train neural network with domain adaptation: the feature generator transforms the source-domain data and the target-domain data into high-level abstract features, making the domains of features difficult to discriminate as far as possible; and the domain discriminator reads the features and tries to discriminate their domains as accurately as possible.

Adversarial samples [28], [29] also contain the adversarial idea. Adversarial samples refer to those samples that have little difference from the real samples, but are classified into a wrong category, or those that have big difference from the real samples, but are classified into a real category with high confidence. In order to learn an object detector that is robust to occlusion and deformation, Wang *et al.* [30] utilize the adversarial learning idea to generate positive examples which are hard for the object detector to recognize.

Adversarial learning, adversarial networks, and adversarial samples all contain adversarial idea, but have different objectives. The already achieved results by applying adversarial idea in AI inspire more researchers to explore GANs.

III. THEORY AND IMPLEMENTATION MODELS OF GANS

A. Basic Theory

The main idea of GAN comes from the Nash equilibrium in game theory [1]. It assumes two game participants: one generator and one discriminator. The generator aims to learn the distribution of real data, while the discriminator aims to correctly determine whether the input data is from the real data or from the generator. In order to win the game, the two participants need to continuously optimize themselves to improve the generation ability and the discrimination ability, respectively. The purpose of the optimization process is to find a Nash equilibrium between the two participants.

The computation procedure and structure of GAN is shown in Fig. 2. Any differentiable function can be used as the generator and the discriminator. Here, we use differentiable functions D and G to represent the discriminator and the generator, and their inputs are real data x and random variables z, respectively. $G(z)$ represents the sample generated by G and obeying the distribution $p_{\rm data}$ of real data. If the input of discriminator D is from the real data x, D should classify it to be true and label it as 1. If the input is from $G(z)$, D should classify it to be false and label it as 0. The purpose of D is to achieve correct classification of the data source, while the purpose of G is to make performance of the generated data $G(z)$ on D (i.e., $D(G(z))$ consistent with the performance of real data x on D (i.e., $D(x)$). The adversarial optimization process improves the performance of D and G gradually. Eventually, when the discrimination ability of D has been improved to a high level but cannot discriminate the data source correctly, it is thought that the generator G has captured the distribution of real data.

B. Learning Method

In this subsection, we discuss the learning and training mechanism of GAN.

First, we describe the optimization of discriminator D given generator G. Similar to the training of Sigmoid function-based classifiers, training the discriminator involves minimizing the cross entropy. The loss function is formulated as below:

$$\begin{aligned} Obj^D(\theta_D,\theta_G) = &-\frac{1}{2}E_{x\sim p_{\rm data}(x)}[\log D(x)] \\ &-\frac{1}{2}E_{z\sim p_z(z)}[\log(1-D(g(z)))] \end{aligned} \tag{1}$$

where x is sampled from real data distribution $p_{\rm data}(x)$, z is sampled from the prior distribution $p_z(z)$ such as uniform or Gaussian distribution, and $E(\cdot)$ represents the expectation. It should be noted that the training data consists of two parts: one part from the real data distribution $p_{\rm data}(x)$ and another part from the generated data distribution $p_g(x)$. This is slightly different from conventional methods for binary classification. Given the generator, we need to minimize (1) to obtain the optimal solution. In continuous space, (1) can be reformulated as below:

$$\begin{aligned} Obj^D(\theta_D, \theta_G) &= -\frac{1}{2}\int_x p_{\mathrm{data}}(x)\log(D(x))dx \\ &\quad -\frac{1}{2}\int_z p_z(z)\log(1-D(g(z)))dz \\ &= -\frac{1}{2}\int_x [p_{\mathrm{data}}(x)\log(D(x)) \\ &\quad + p_g(x)\log(1-D(x))]dx. \end{aligned} \tag{2}$$

For any $(m, n) \in \mathbb{R}^2 \setminus \{0, 0\}$ and $y \in [0, 1]$, the expression

$$-m\log(y) - n\log(1-y) \tag{3}$$

achieves its minimum value at $y = m/(m+n)$. Hence, given generator G, the objective function (2) achieves its minimum value at

$$D_G^*(x) = \frac{p_{\mathrm{data}}(x)}{p_{\mathrm{data}}(x) + p_g(x)}. \tag{4}$$

This is the optimal solution of discriminator D. From (4), the discriminator of GAN estimates the ratio of two probability densities, which is the key difference from Markov chain or lower bound based methods.

On the other hand, $D(x)$ denotes the probability of x sampled from the real data rather than the generated data. If the input data is from real data x, the discriminator strives to make $D(x)$ approach 1. If the input data is from the generated data $G(z)$, the discriminator strives to make $D(G(z))$ approach 0 while the generator G tries to make it approach 1. Since this is a zero-sum game between G and D, the loss function of G is $Obj^G(\theta_G) = -Obj^D(\theta_D, \theta_G)$. Therefore, the optimization of GAN can be formulated as a minimax problem:

$$\begin{aligned} \min_G \max_D \{f(D, G) = nE_{x\sim p_{\mathrm{data}}(x)}[\log D(x)] \\ + E_{z\sim p_z(z)}[\log(1 - D(G(z)))]\}. \end{aligned} \tag{5}$$

In summary, for learning the parameters of GAN, we need to train the discriminator D to maximize the accuracy of discriminating the input data from the real data x or the generated data $G(z)$. In addition, we need to train the generator G to minimize $\log(1 - D(G(z)))$. Here we can use an alternative training method. First, we fix G and optimize D to maximize the discrimination accuracy of D. Then, we fix D and optimize G to minimize the discrimination accuracy of D. This process alternates and we could achieve the global optimal solution if and only if $p_{\mathrm{data}} = p_g$. In the training process, we empirically update the parameters of D for k times and then update the parameters of G once.

C. GAN Variants

Since Goodfellow's proposal of GAN [1], a number of GAN variants have been created. The major innovative points include model structure improvement, theoretical extension, novel applications, etc. The computation procedures and structures of some GAN variants are shown in Fig. 3.

Since the real data and generated data can have very little overlap, the Jensen-Shannon divergence of objective function can be a constant, which causes the vanishing gradient problem while using the gradient descent method to train GANs. To tackle the vanishing gradient problem, Arjovsky *et al.* [31] propose Wasserstein GAN (WGAN) by using the Earth-Mover distance to replace the Jensen-Shannon divergence for evaluating the distribution distance between real data and the generated data. They use a critic function f that builds on Lipschitz constraint to represent the discriminator. WGAN makes significant progress towards stable training of GANs, but can still generate low-quality samples or fail to converge in some settings. In light of that, Gulrajani *et al.* [39] find that the training failures are often due to the use of weight clipping in WGAN to enforce a Lipschitz constraint on the critic, which can lead to pathological behavior. They propose an alternative method for enforcing the Lipschitz constraint: instead of clipping weights, penalize the norm of the gradient of the critic with respect to its input. Their method converges faster and generates higher-quality samples than WGAN with weight clipping.

Another issue about GAN is that the discriminator has an infinite modeling ability and can distinguish between real samples and generated samples regardless of their complexity, which easily causes over-fitting. To limit the modeling ability of the discriminator, Qi [32] propose Loss-Sensitive GAN (LS-GAN), which demands the objective function to satisfy the Lipschitz constraint. In addition, they give some quantitative results when the gradients vanish. It should be pointed out that WGAN and LS-GAN do not change the GAN structure, but improve the parameter learning and optimization method.

In general, only the label about data source is required for training GANs. Odena [33] proposes Semi-GAN by adding labels of real data to the training of discriminator D. Furthermore, Mirza *et al.* [34] propose to add auxiliary information y to G, D, and real data x for GAN modeling. Here y can be labels or other auxiliary information. Conventional GANs aim to learn a generative model to map the latent variable distribution to complex real data distribution. Donahue *et al.* [35] propose Bidirectional GANs (BiGANs) to map the real data to the latent variable space, thereby achieving feature learning. On top of the basic structure of GAN, BiGANs add an extra decoder Q to map the real data x to latent space, so that the optimization problem is converted to $\min_{G,Q}\max_D f(D, Q, G)$.

Conventional GANs can learn some semantic features, but cannot capture the relationship between the dimension of random variables z and specific semantics. Chen *et al.* [36] propose InfoGAN to capture the mutual information between a small subset of latent variables and the observation. In particular, the input noise vector of G is decomposed into two parts: z and c. z is treated as source of incompressible noise and c is called the latent code to represent the structured semantic data distribution. Conventional GANs set $p_G(x) = p_G(x|c)$, but in fact c and the output of generator G are strongly correlated. Let $G(z, c)$ denote the output of generator G. InfoGAN uses the mutual information $I(c; G(z, c))$ to represent the correlation level of two samples. Their objective function is formulated as below:

$$\min_G \max_D \{f_I(D, G) = f(D, G) - \lambda I(c; G(z, c))\}. \tag{6}$$

Since the posterior probability $p(c|x)$ cannot be computed explicitly, its lower bound can be estimated via variational information maximization.

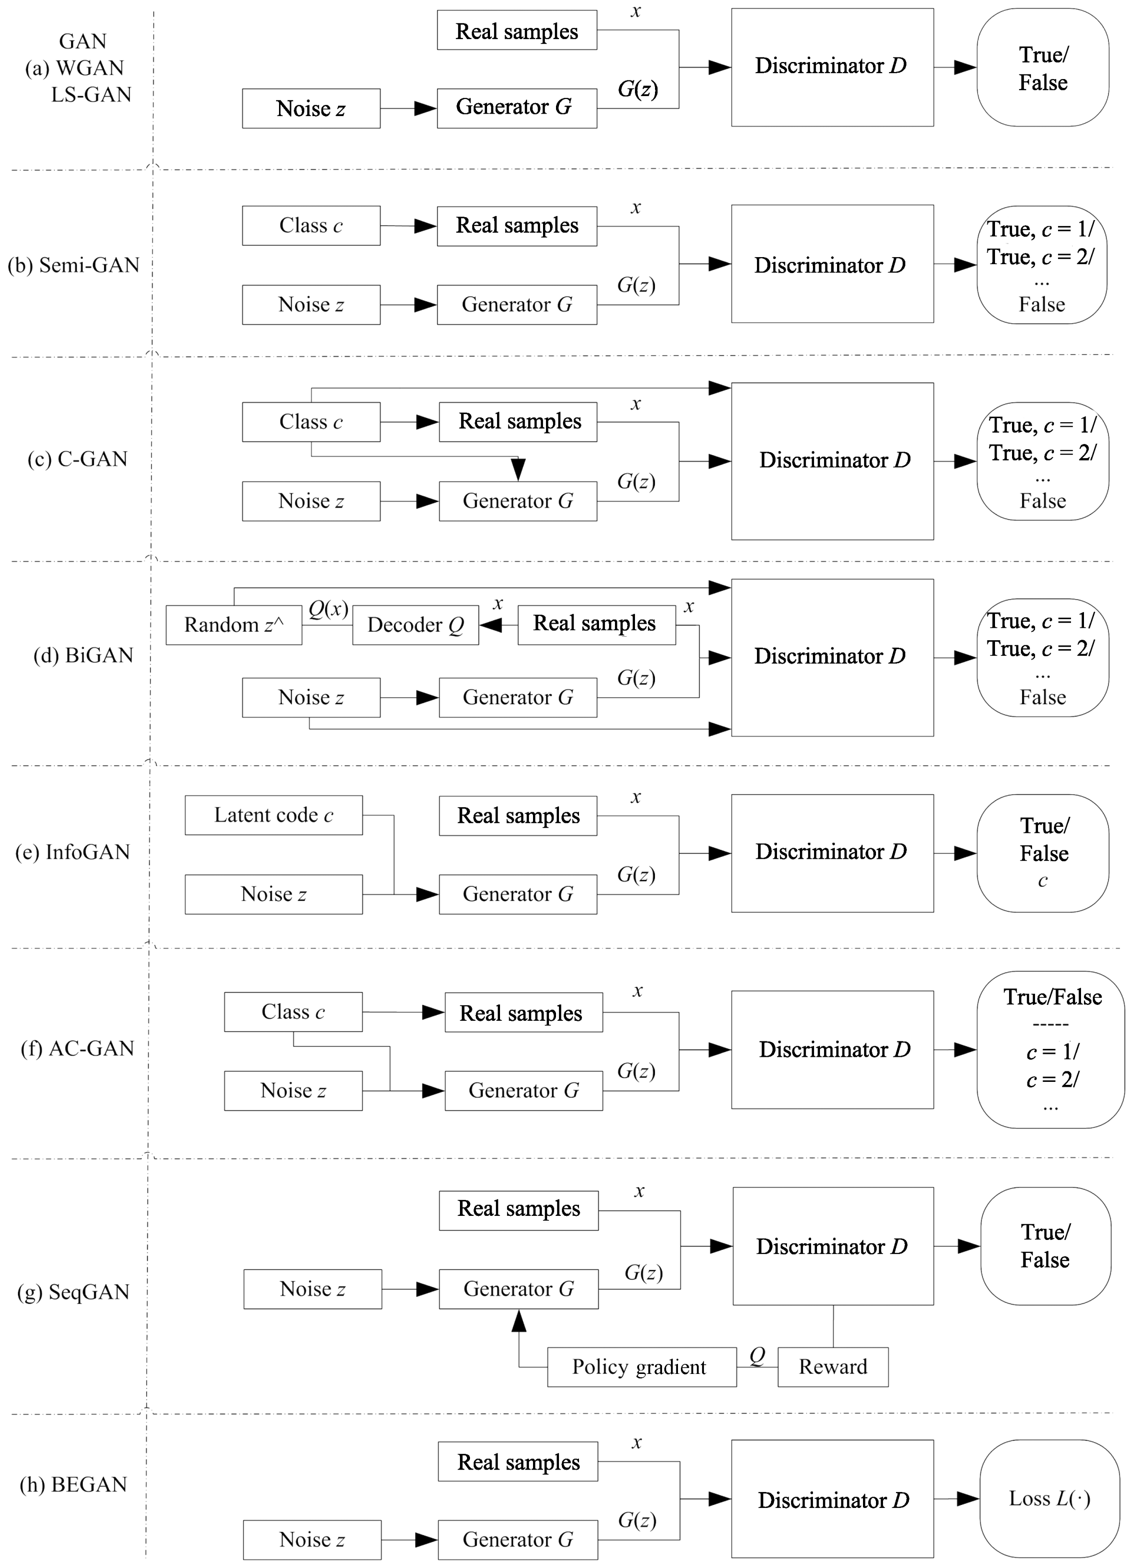

Fig. 3. Computation procedures and structures of some GAN variants: (a) GAN [1], WGAN [31], LS-GAN [32], (b) Semi-GAN [33], (c) C-GAN [34], (d) BiGAN [35], (e) InfoGAN [36], (f) AC-GAN [37], (g) SeqGAN [7], and (h) BEGAN [38].

Odena *et al.* [37] propose auxiliary classifier GAN (AC-GAN) for semi-supervised synthesis. Their objective function consists of two parts: the log-likelihood of the correct data source and that of the correct class. The key of AC-GAN is that it can incorporate label information into the generator and adjust the objective function for the discriminator. In consequence, the generation and discrimination abilities of GAN are improved.

Yu *et al.* [7] propose SeqGAN to generate data sequences. It is the first work that extends GANs for generating sequences of discrete tokens. SeqGAN models G as a stochastic policy in reinforcement learning to overcome the generator differentiation problem. It uses the policy gradient reinforcement learning to backpropagate the error from D. Experiments show that SeqGAN can achieve preferable results on speech, poem, and music generation.

Berthelot *et al.* [38] propose BEGAN, which is a new equilibrium enforcing method paired with a loss derived from the Wasserstein distance for training auto-encoder based GANs. While typical GANs try to match data distributions directly, BEGAN matches auto-encoder loss distributions. This method balances the generator and discriminator during training, and provides a new approximate convergence measure. The authors also derive a way of controlling the trade-off between image diversity and visual quality. On the image generation task, they set a new milestone in visual quality.

IV. Applications of GANs

GANs can be used to generate samples with the same distribution as real data, e.g., generating photorealistic images. GANs can also be used to tackle the problem of insufficient training samples for supervised or semi-supervised learning. In addition, GANs have been applied for speech and language processing, such as generating dialogues. In this section, we discuss the application range of GANs.

A. Applications to Image and Vision Computing

GANs can generate image samples with the same distribution of real images. One typical application is from [40], where Ledig *et al.* present SRGAN for image super-resolution. They use VGG network as the discriminator and residual network as the generator. Experimental results show that SRGAN can get rich texture details for the estimation of photorealistic super-resolution images.

BEGAN [38] is able to generate high-quality face samples at resolutions of 128×128. Varied poses, expressions, genders, skin colors, light exposure, and facial hair are observed from the generated samples, as shown in Fig. 4.

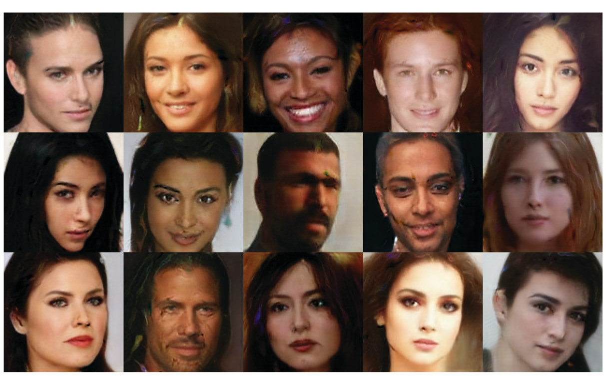

Fig. 4. Face samples generated by BEGAN [38].

GANs can be used to generate driving scenarios. Santana *et al.* [41] propose to generate images with the same distribution as real driving scenarios. Their driving simulator model is an autoencoder trained with generative adversarial network based costs coupled with a recurrent neural network transition model. Results show that GANs can generate realistic looking images of the road, and thus can be applied in autonomous driving for unsupervised or semi-supervised learning.

Gou *et al.* [42], [43] propose to learn from both real images and synthetic images for accurate eye detection. But both the synthetic and real images they use have limitations, because the synthetic images do not contain eyes with glasses while the real images do not cover diverse illuminations and appearance variations. Shrivastava *et al.* [44] propose SimGAN to learn from simulated and unsupervised images (S + U learning) for bridging the gap between synthetic and real image datasets. The framework of SimGAN is shown in Fig. 5. They learn the generator to refine synthetic data so that it follows the distribution of real data while maintaining the annotations of synthetic data.

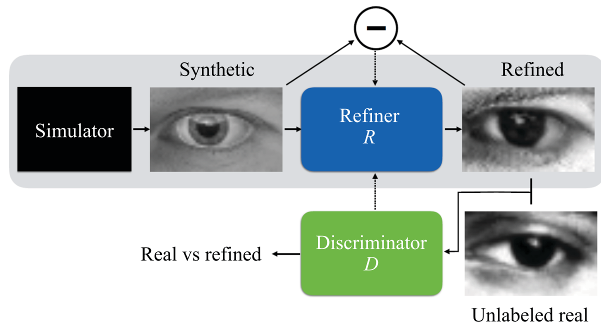

Fig. 5. Framework of SimGAN [44].

Image-to-image translation is a class of vision applications in which the goal is to learn the mapping between an input image and an output image. Huang *et al.* [45] propose Two-Pathway GAN (TP-GAN) for photorealistic frontal view synthesis from a single face image, by simultaneously perceiving global structures and local details. The synthesized identity preserving image can be used for downstream tasks like face recognition. The training phase of TP-GAN requires paired examples of identity preserving frontal view image and face image under a different pose. Nevertheless, paired training data are not available for many tasks. In light of that, Zhu *et al.* [46] propose CycleGAN for learning to translate an image from a source domain to a target domain in the absence of paired examples. Their method is a general-purpose one and can be applied to a wide range of image-to-image translation tasks, including style transfer, object transfiguration, attribute transfer, and photo enhancement, as shown in Fig. 6.

B. Applications to Speech and Language Processing

Recently, there are some GANs based applications for speech and language processing. Li *et al.* [6] use GANs to capture the relevance of dialogue and generate corresponding text. Zhang *et al.* [47] propose to generate realistic sentence with GANs, by using long short-term memory and convolutional neural networks for adversarial training. SeqGAN [7] employs reinforcement learning to generate speech and language, poem, and music. Pascual *et al.* [48] propose the use of GANs for speech enhancement, called SEGAN. They operate at the waveform level, train the model end-to-end, and incorporate 28 speakers and 40 different noise conditions into the same model, such that model parameters are shared. They evaluate the proposed model using an independent, unseen test set with two speakers and 20 alternative noise conditions. The

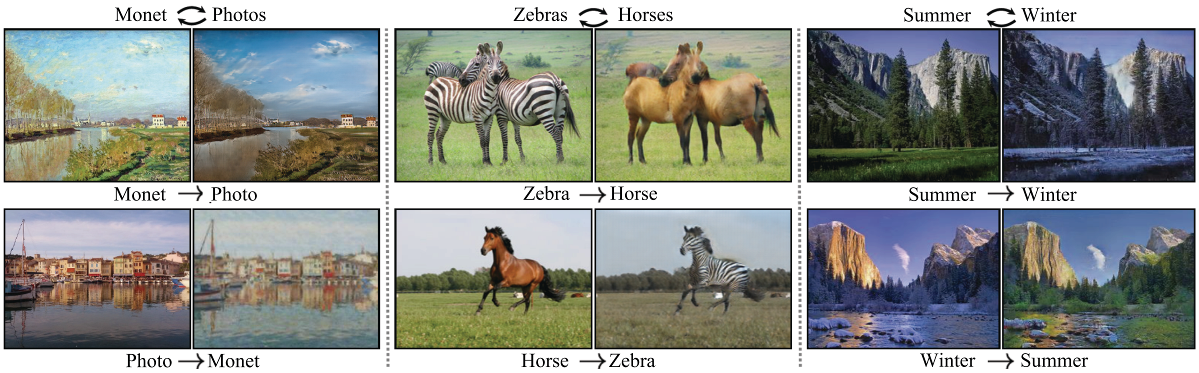

Fig. 6. Example results of CycleGAN for unpaired image-to-image translation [46].

enhanced samples confirm the viability of SEGAN.

Some researchers propose to synthesize images from text descriptions. Reed *et al.* [49] present a GANs-based method for generating images from the text. Text encoding is used by both the generator and the discriminator. Experimental results show that it can get preferable images from text descriptions. However, their generated samples fail to contain necessary details and vivid object parts. Zhang *et al.* [50] propose StackGAN for text to photorealistic image synthesis. Given text descriptions, Stage-I of StackGAN sketches rough shapes and basic colors of objects, yielding low resolution images. Stage-II of StackGAN takes Stage-I results and text descriptions as inputs, and generates high resolution images with photorealistic details. StackGAN can generate realistic 256×256 images conditioned on only text descriptions. Fig. 7 shows some example images generated by StackGAN from unseen text descriptions.

C. Other Applications

GANs can connect with reinforcement learning, such as in the aforementioned SeqGAN [7]. Some researchers also connect GANs with imitation learning [51], [52] or Actor-critic based methods [53]. Next we take three examples to indicate the usefulness of GANs in a range of applications.

Hu *et al.* [8] propose a GAN-based algorithm named MalGAN to generate adversarial malware examples, which are able to bypass black-box machine learning based detection models. MalGAN uses a substitute detector to fit the black-box malware detection system. A generative network is trained to minimize the generated adversarial examples' malicious probabilities predicted by the substitute detector. MalGAN is able to decrease the malware detection rate to nearly zero and make the retraining based defensive method against adversarial examples hard to work. This is an important conclusion because malware detection algorithms cannot be used in real-world applications if they are easy to be bypassed by adversarial examples.

Chidambaram *et al.* [9] present a general framework for style transfer, which they term style transfer generative adversarial networks (STGANs) as an extension of GANs. They use a discriminator to regularize a generator with an otherwise separate loss function. Their approach is applied to the task of learning to play chess in the style of a specific player, and produce empirical evidence for the viability of STGANs.

Choi *et al.* [54] propose medical generative adversarial network (medGAN) to generate realistic synthetic electronic health records (EHRs). Based on an input EHR dataset, medGAN can generate high-dimensional discrete variables (e.g., binary and count features) via a combination of an autoencoder and GANs. To demonstrate feasibility, they showed that medGAN generates synthetic EHR datasets that achieve comparable performance to real data on many experiments including distribution statistics, predictive modeling tasks and medical expert review.

V. IMPACT AND PROSPECT

A. Significance and Advantages

GANs have great significance to the development of generative models. As a powerful class of generative methods, GANs solve the problem of generating data that can be naturally interpreted. Especially for the generation of high-dimensional data, the adopted neural network structure does not limit the generation dimension, which greatly broadens the scope of the generated data samples. Besides, the neural network structure can integrate various loss functions, thereby increasing the degree of freedom of the model design. In general, the training process of GANs uses two adversarial neural networks as training criterion and can be trained by backpropagation. The training does not rely on the inefficient Markov chain method, nor approximate inference. There is no complex variational lower bound, which greatly reduces the training difficulty and improves the training efficiency. The generation process of GANs does not require tedious sampling sequence, but can directly sample and predict new samples, which improve the efficiency of generating new samples. The adversarial training discards direct replication or average of real data, thereby increasing the diversity of the generated samples. In practice, the samples generated by GANs are easy to understand for humans. For example, GANs can generate very sharp and realistic images. In brief, GANs provide a promising solution for creatively generating data that are meaningful to humans.

Not only have GANs made great contributions to the development of generative models, but they are also meaningful and instructive for semi-supervised learning. As we know, the learning process of GANs does not require data labels except the data source. Although the objective of GANs is not semi-supervised learning, the training process of GANs can be used to achieve pre-training using unlabeled data. For example, we

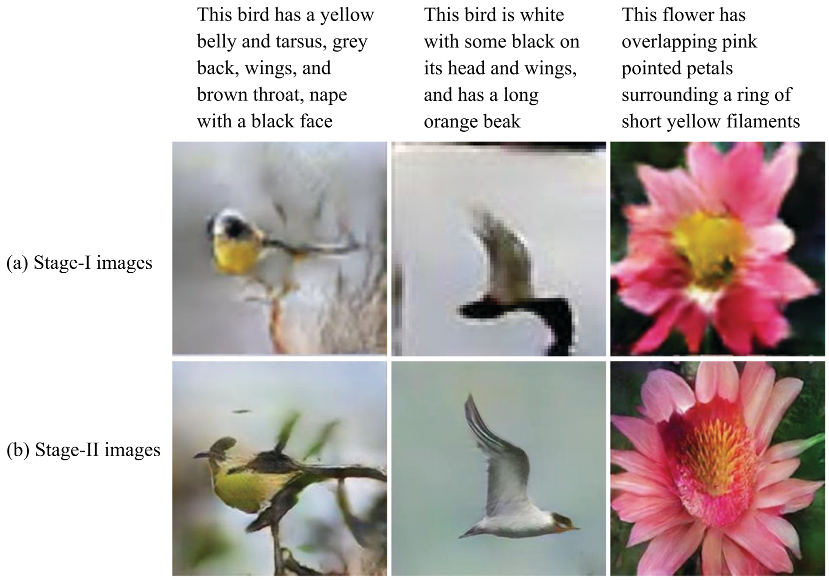

Fig. 7. Photorealistic images generated by StackGAN from text descriptions [50].

can first use a large amount of unlabeled data to train GANs; based on the understanding of the trained GANs over the unlabeled data, we can then use a small amount of labeled data to train the discriminative model for classification and regression tasks.

B. Limitations and Development Trends

GANs have solved a lot of problems for generative models and brought inspiration to other AI methods, but they still have limitations. GANs adopt the adversarial learning idea, but convergence of the model and existence of equilibrium point have not been proved yet. The training process needs to ensure balance and synchronization of two adversarial networks, otherwise it is difficult to obtain good training results. However, it is difficult to control the synchronization of the two adversarial networks, so the training process may be unstable. In addition, as generative models based on neural networks, GANs have the common defect (i.e., poor interpretability) of neural networks. Furthermore, although the samples generated by GANs are diverse, there exists the collapse mode problem [5]. Mode collapse refers to scenarios in which the generator makes multiple images that contain the same color or texture themes, thereby having little difference for human understanding.

Although GANs have some limitations, it is undeniable that the research progress of GANs has revealed their broad prospects. New techniques dedicated to reducing the limitations are continually emerging. For example, Wasserstein GAN [31], [39] greatly overcomes the training instability problem, and partially solves the collapse mode problem at the same time. How to completely avoid collapse mode and further optimize the training process remains a research direction of GANs. Furthermore, the theory about model convergence and the existence of equilibrium point remain important research subjects in the near future.

The above research directions focus on better solving the drawbacks of GANs. From the perspective of developing and applying GANs, how to generate a variety of data that can interact with humans from simple random inputs is an important research direction in the near future. From the perspective of combining GANs and other methods, how to integrate GANs with feature learning, imitation learning, and reinforcement learning to develop new AI applications and promote the development of these methods is very meaningful. In the long run, how to use GANs to promote the development and application of AI, enhance the ability of AI to understand the world, and even stimulate the creativity of AI are important problems that should be considered by researchers.

C. Relation Between GANs and Parallel Intelligence

In 2004, Fei-Yue Wang [55], [56] proposed the parallel systems theory and ACP (artificial societies, computational experiments, and parallel execution) approach for modeling and control of complex systems. Parallel systems emphasize virtual-real interaction. Artificial systems are constructed to represent the actual system, computational experiments are utilized to learn and evaluate various computational models, and parallel execution is implemented to improve the performance of the actual system. In parallel systems, the artificial systems and the actual system work together in a virtual-real interactive manner [57], [58]. The parallel systems theory and ACP approach have now evolved into a more generalized parallel intelligence theory [59]. In the training process of GANs, the real data samples and the generated data samples interact with each other via the adversarial networks, and the trained generator can generate more virtual samples than the real samples. GANs can deepen the philosophy of parallel systems' virtual-real interaction and integration. As a powerful class of effective generative models, GANs can merge into the systematic research of parallel intelligence. In this subsection, we discuss the relation between GANs and parallel intelligence from three aspects.

1) GANs and Parallel Vision: Parallel vision [60] is an extension of ACP approach into the vision computing field. Fig. 8

shows the basic framework and architecture of parallel vision. Integrating multiple technologies (such as computer graphics, virtual reality, machine learning, and knowledge automation) and utilizing a new ACP (artificial scenes, computational experiments, and parallel execution) approach, parallel vision aims to establish systematic theories and methods for visual perception and understanding of complex scenes. Parallel vision first constructs artificial scenes to simulate and represent complex real scenes, making it possible to collect and annotate large-scale diversified image datasets. Computational experiments are then utilized to design and evaluate a variety of vision algorithms. Finally, parallel execution is used to online optimize the vision system. In parallel vision, the generation of artificial-scene images can be realized by GANs, as shown in Fig. 4 and Fig. 6. GANs can generate large-scale diversified image datasets, which can be combined with real datasets to train the vision models. This helps improve the generalization ability of vision models.

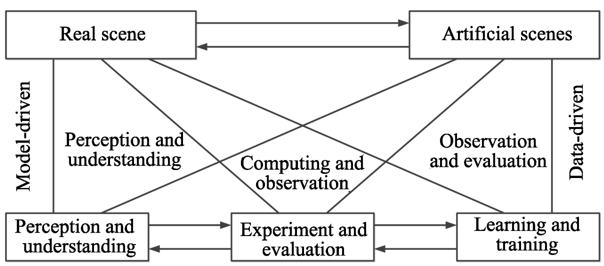

Fig. 8. Basic framework and architecture for parallel vision.

2) GANs and Parallel Control: Parallel control [61]−[63] is a specific application of ACP approach in the field of complex system control. Fig. 9 shows the structure of parallel control systems. In parallel control, artificial systems are used for modeling and representation, computational experiments are utilized for analysis and evaluation, and parallel execution is conducted for control and management of complex systems. Parallel control can be considered as the extension of feedback control, especially adaptive control, for dealing with problems involved with both engineering and social complexities. In addition to the generation of artificial systems and the analysis of computational experiments, the parallel execution between artificial systems and the actual system can also be simulated using GANs. On one hand, GANs can be used to conduct predictive learning of the artificial systems and feedback learning of the actual system. On the other hand, GANs can be used to realize imitation learning and reinforcement learning of the control unit.

3) GANs and Parallel Learning: Parallel learning [64] is the extension of ACP theory in the learning field. Fig. 10 shows the theoretical framework of parallel learning. The emphases of parallel learning include: using predictive learning to solve the problem of exploring time-varying data, using ensemble learning to solve the problem of exploring spatial-distributed data, and using prescriptive learning to solve the problem of exploring the direction of data generation. Parallel learning can be considered as a new theoretical framework for machine learning, and is closely related to parallel vision and parallel control. GANs can be combined with parallel learning in terms of big data generation, computational experiments based predictive learning, and so on.

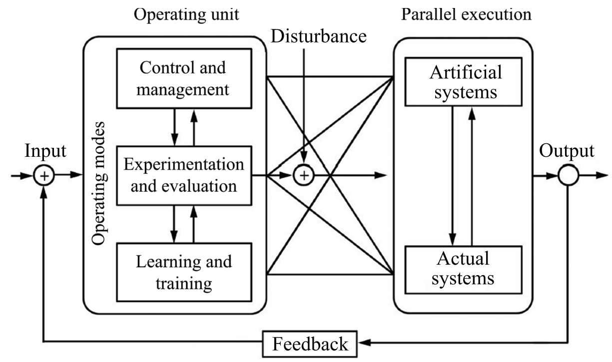

Fig. 9. Structure of parallel control systems.

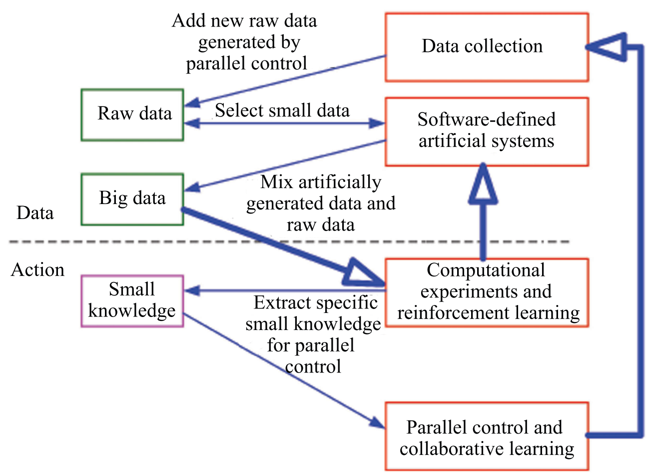

Fig. 10. Theoretical framework of parallel learning.

VI. Concluding Remarks

In this paper, we survey the state of the art of GANs. Since the proposal of GAN in 2014 by Goodfellow, this model has been receiving increasing attention from the AI community. The core idea of GANs originates from two-player zero-sum game in game theory. A GAN usually comprises a generator and a discriminator, which are trained iteratively in an adversarial learning manner, approaching Nash equilibrium. As a powerful class of generative models, GANs do not estimate the distribution of data samples explicitly, but learn to generate new samples that conform to the same distribution as the real samples. The ability to generate "infinite" new samples from potential distribution has great application value in many fields such as image and vision computing, speech and language processing, and information security.

We also investigate the development trends of GANs, and discuss the relation between GANs and parallel intelligence in particular. In our opinion, GANs can deepen the philosophy of parallel systems' virtual-real interaction and integration, and provide specific and substantial algorithmic support for the ACP approach. In parallel vision, parallel control, and parallel learning systems, GANs can learn to generate data samples that have the same distribution as the real samples, thereby supporting the research and application of parallel systems. In summary, GANs as powerful generative models can merge into the systematic study of parallel intelligence.

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 27*, Montreal, Quebec, Canada, 2014, pp. 2672−2680.

[2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. New York, USA: MIT Press, 2016.

[3] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," arXiv: 1511.06434, 2015.

[4] L. J. Ratliff, S. A. Burden, and S. S. Sastry, "Characterization and computation of local Nash equilibria in continuous games," in *Proc. 51st Annu. Allerton Conf. Communication, Control, and Computing (Allerton)*, Monticello, IL, USA, 2013, pp. 917−924.

[5] I. Goodfellow, "NIPS 2016 tutorial: generative adversarial networks," arXiv: 1701.00160, 2016.

[6] J. W. Li, W. Monroe, T. L. Shi, S. Jean, A. Ritter, and D. Jurafsky, "Adversarial learning for neural dialogue generation," arXiv: 1701.06547, 2017.

[7] L. T. Yu, W. N. Zhang, J. Wang, and Y. Yu, "SeqGAN: sequence generative adversarial nets with policy gradient," arXiv: 1609.05473, 2016.

[8] W. W. Hu and Y. Tan, "Generating adversarial malware examples for black-box attacks based on GAN," arXiv: 1702.05983, 2017.

[9] M. Chidambaram and Y. J. Qi, "Style transfer generative adversarial networks: Learning to play chess differently," arXiv: 1702.06762, 2017.

[10] Y. Bengio, "Learning deep architectures for AI," *Found. Trends Mach. Learn.*, vol. 2, no. 1, pp. 1−127, Jan. 2009.

[11] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," arXiv: 1312.6114, 2013.

[12] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," arXiv: 1401. 4082, 2014.

[13] G. E. Hinton, T. J. Sejnowski, and D. H. Ackley, "Boltzmann machines: constraint satisfaction networks that learn," Carnegiemellon Univ., Pittsburgh, PA, USA, Tech. Rep. CMU-CS-84-119, 1984.

[14] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for Boltzmann machines," *Cognit. Sci.*, vol. 9, no. 1, pp. 147−169, Jan. 1985.

[15] G. E. Hinton, S. Osindero, and Y. W. Teh, "A fast learning algorithm for deep belief nets," *Neural Computat.*, vol. 18, no. 7, pp. 1527−1554, May 2006.

[16] Y. Bengio, É. Thibodeau-Laufer, G. Alain, and J. Yosinski, "Deep generative stochastic networks trainable by backprop," arXiv: 1306.1091, 2013.

[17] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504−507, Jul. 2006.

[18] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436−444, May 2015.

[19] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. 25th Int. Conf. Neural Information Processing Systems*, Lake Tahoe, Nevada, USA, 2012, pp. 1097−1105.

[20] K. M. He, X. Y. Zhang, S. Q. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Las Vegas, NV, USA, 2016, pp. 770−778.

[21] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82−97, Nov. 2012.

[22] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems 27: Annu. Conf. Neural Information Processing Systems 2014*, Montreal, Quebec, Canada, 2014, pp. 3104−3112.

[23] D. He, W. Chen, L. W. Wang, and T. Y. Liu, "A game-theoretic machine learning approach for revenue maximization in sponsored search," arXiv: 1406.0728, 2014.

[24] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484−489, Jan. 2016.

[25] J. Schmidhuber, "Learning factorial codes by predictability minimization," *Neural Computat.*, vol. 4, no. 6, pp. 863−879, Nov. 1992.

[26] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096−2030, Jan. 2016.

[27] W. Z. Chen, H. Wang, Y. Y. Li, H. Su, Z. H. Wang, C. H. Tu, D. Lischinski, D. Cohen-Or, and B. Q. Chen, "Synthesizing training images for boosting human 3D pose estimation," in *Proc. 4th Int. Conf. 3D Vision (3DV)*, Stanford, CA, USA, 2016, pp. 479−488.

[28] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," arXiv: 1312. 6199, 2014.

[29] P. McDaniel, N. Papernot, and Z. B. Celik, "Machine learning in adversarial settings," *IEEE Secur. Priv.*, vol. 14, no. 3, pp. 68−72, May−Jun. 2016.

[30] X. L. Wang, A. Shrivastava, and A. Gupta, "A-Fast-RCNN: hard positive generation via adversary for object detection," arXiv: 1704.03414, 2017.

[31] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," arXiv: 1701.07875, 2017.

[32] G. J. Qi, "Loss-sensitive generative adversarial networks on Lipschitz densities," arXiv: 1701.06264, 2017.

[33] A. Odena, "Semi-supervised learning with generative adversarial networks," arXiv: 1606.01583, 2016.

[34] M. Mirza and S. Osindero, "Conditional generative adversarial nets," arXiv: 1411.1784, 2014.

[35] J. Donahue, P. Krähenbühl, and T. Darrell, "Adversarial feature learning," arXiv: 1605.09782, 2017.

[36] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel, "InfoGAN: interpretable representation learning by information maximizing generative adversarial nets," in *Proc. 30th Conf. Neural Information Processing Systems*, Barcelona, Spain, 2016, pp. 2172−2180.

[37] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier GANs," arXiv: 1610.09585, 2017.

[38] D. Berthelot, T. Schumm, and L. Metz, "BEGAN: boundary equilibrium generative adversarial networks," arXiv: 1703.10717, 2017.

[39] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of Wasserstein GANs," arXiv: 1704.00028, 2017.

[40] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. H. Wang, and W. Z. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," arXiv: 1609.04802, 2017.

[41] E. Santana and G. Hotz, "Learning a driving simulator," arXiv: 1608. 01230, 2016.

[42] C. Gou, Y. Wu, K. Wang, F. -Y. Wang, and Q. Ji, "Learning-by-synthesis for accurate eye detection," in *Proc. 23rd Int. Conf. Pattern Recognition (ICPR)*, Cancun, Mexico, 2016, pp. 3362−3367.

[43] C. Gou, Y. Wu, K. Wang, K. F. Wang, F. -Y. Wang, and Q. Ji, "A joint cascaded framework for simultaneous eye detection and eye state estimation," *Pattern Recognit.*, vol. 67, pp. 23−31, Jul. 2017.

[44] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. D. Wang, and R. Webb, "Learning from simulated and unsupervised images through adversarial training," arXiv: 1612.07828, 2016.

[45] R. Huang, S. Zhang, T. Y. Li, and R. He, "Beyond face rotation: global and local perception GAN for photorealistic and identity preserving frontal view synthesis," arXiv: 1704.04086, 2017.

[46] J. Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," arXiv: 1703. 10593, 2017.

[47] Y. Z. Zhang, Z. Gan, and L. Carin, "Generating text via adversarial training," *Proc. Workshop on Adversarial Training*, Barcelona, Spain, 2016.

[48] S. Pascual, A. Bonafonte, and J. Serrá, "SEGAN: speech enhancement generative adversarial network," arXiv: 1703.09452, 2017.

[49] S. Reed, Z. Akata, X. C. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," in *Proc. 33rd Int. Conf. Machine Learning*, New York, NY, USA, vol. 48, pp. 1060−1069, 2016.

[50] H. Zhang, T. Xu, H. S. Li, S. T. Zhang, X. L. Huang, X. G. Wang, and D. Metaxas, "StackGAN: text to photo-realistic image synthesis with stacked generative adversarial networks," arXiv: 1612.03242, 2016.

[51] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Proc. 30th Conf. Neural Information Processing Systems*, Barcelona, Spain, 2016, pp. 4565−4573.

[52] C. Finn, P. Christiano, P. Abbeel, and S. Levine, "A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models," arXiv: 1611.03852, 2016.

[53] D. Pfau and O. Vinyals, "Connecting generative adversarial networks and actor-critic methods," arXiv: 1610.01945, 2017.

[54] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. M. Sun, "Generating multi-label discrete electronic health records using generative adversarial networks," arXiv: 1703.06490, 2017.

[55] F. -Y. Wang, "Parallel system methods for management and control of complex systems," *Control Dec.*, vol. 19, no. 5, pp. 485−489, 514, May 2004.

[56] F. -Y. Wang, "Computational experiments for behavior analysis and decision evaluation of complex systems," *J. Syst. Simulat.* vol. 16, no. 5, pp. 893−897, May 2004.

[57] F. -Y. Wang, J. Zhang, Q. L. Wei, X. H. Zheng, and L. Li, "PDP: parallel dynamic programming," *IEEE/CAA J. Automat. Sinica*, vol. 4, no. 1, pp. 1−5, Jan. 2017.

[58] T. X. Bai, S. Wang, Z. Shen, D. P. Cao, N. N. Zheng, and F. -Y. Wang, "Parallel robotics and parallel unmanned systems: framework, structure, process, platform and applications," *Acta Automat. Sinica*, vol. 43, no. 2, pp. 161−175, Feb. 2017.

[59] F. -Y. Wang, X. Wang, L. X. Li, and L. Li, "Steps toward parallel intelligence," *IEEE/CAA J. Automat. Sinica*, vol. 3, no. 4, pp. 345−348, Oct. 2016.

[60] K. F. Wang, C. Gou, and F. -Y. Wang, "Parallel vision: an ACP-based approach to intelligent vision computing," *Acta Automat. Sinica*, vol. 42, no. 10, pp. 1490−1500, Oct. 2016.

[61] F. -Y. Wang, "On the modeling, analysis, control and management of complex systems," *Complex Syst. Complex. Sci.*, vol. 3, no. 2, pp. 26−34, Jun. 2006.

[62] F. -Y. Wang, D. R. Liu, G. Xiong, C. J. Cheng, and D. B. Zhao, "Parallel control theory of complex systems and applications," *Complex Syst. Complex. Sci.*, vol. 9, no. 3, pp. 1−12, Sep. 2012.

[63] F. -Y. Wang, "Parallel control: a method for data-driven and computational control," *Acta Automat. Sinica*, vol. 39, no. 4, pp. 293−302, Apr. 2013.

[64] L. Li, Y. L. Lin, D. P. Cao, N. N. Zheng, and F. -Y. Wang, "Parallel learning — a new framework for machine learning," *Acta Automat. Sinica*, vol. 43, no. 1, pp. 1−8, Jan. 2017.

Kunfeng Wang (M'17) received the Ph.D. degree in control theory and control engineering from the Graduate University of Chinese Academy of Sciences, Beijing, China, in 2008. He joined the Institute of Automation, Chinese Academy of Sciences and became an associate professor at the State Key Laboratory of Management and Control for Complex Systems. From December 2015 to January 2017, he was a visiting scholar at the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA. His research interests include intelligent transportation systems, intelligent vision computing, and machine learning.

Chao Gou received the bachelor degree in automation from the University of Electronic Science and Technology of China, Chengdu, China, in 2012. He is currently working toward the Ph.D. degree in control theory and control engineering at the University of Chinese Academy of Sciences, Beijing, China. His research interests include intelligent transportation systems, image processing, and pattern recognition.

Yanjie Duan received the bachelor degree in automation from Northwestern Polytechnical University, Xi'an, China, in 2012. She is currently working toward the Ph.D. degree in control theory and control engineering at the University of Chinese Academy of Sciences, Beijing, China. Her research interests include intelligent transportation systems, machine learning and its application.

Yilun Lin is currently working toward the Ph.D. degree in control theory and control engineering at the University of Chinese Academy of Sciences, Beijing, China. His research interests include social computing, intelligent transportation systems, deep learning and reinforcement learning.

Xinhu Zheng received the B.S. degree in control science and engineering from Zhejiang University, Hangzhou, China, in 2011. He is currently working toward the Ph.D. degree in computer science and engineering at the University of Minnesota, Minneapolis, MN, USA. His research interests include social computing, machine learning, and data analytics.

Fei-Yue Wang (S'87-M'89-SM'94-F'03) received the Ph.D. degree in computer and systems engineering from Rensselaer Polytechnic Institute, Troy, New York in 1990. He joined the University of Arizona in 1990 and became a professor and Director of the Robotics and Automation Laboratory (RAL) and Program in Advanced Research for Complex Systems (PARCS). In 1999, he founded the Intelligent Control and Systems Engineering Center at the Institute of Automation, Chinese Academy of Sciences (CAS), Beijing, China, under the support of the Outstanding Oversea Chinese Talents Program from the State Planning Council and "100 Talent Program" from CAS, and in 2002, was appointed as the Director of the Key Laboratory of Complex Systems and Intelligence Science, CAS. In 2011, he became the State Specially Appointed Expert and the Director of the State Key Laboratory of Management and Control for Complex Systems. Dr. Wang's current research focuses on methods and applications for parallel systems, social computing, and knowledge automation. He was the Founding Editor-in-Chief of the *International Journal of Intelligent Control and Systems* (1995−2000), Founding EiC of *IEEE ITS Magazine* (2006−2007), EiC of *IEEE Intelligent Systems* (2009−2012), and EiC of *IEEE Transactions on ITS* (2009−2016). Currently he is EiC of *China's Journal of Command and Control*. Since 1997, he has served as General or Program Chair of more than 20 IEEE, INFORMS, ACM, ASME conferences. He was the President of IEEE ITS Society (2005−2007), Chinese Association for Science and Technology (CAST, USA) in 2005, the American Zhu Kezhen Education Foundation (2007−2008), and the Vice President of the ACM China Council (2010−2011). Since 2008, he is the Vice President and Secretary General of Chinese Association of Automation. Dr. Wang is elected Fellow of IEEE, INCOSE, IFAC, ASME, and AAAS. In 2007, he received the 2nd Class National Prize in Natural Sciences of China and awarded the Outstanding Scientist by ACM for his work in intelligent control and social computing. He received IEEE ITS Outstanding Application and Research Awards in 2009 and 2011, and IEEE SMC Norbert Wiener Award in 2014.