

Mitigating Mode Collapse using Multiple GANs Training System

Joo Yong Shim[†] · Jean Seong Bjorn Choe^{††} · Jong-Kook Kim^{†††}

ABSTRACT

Generative Adversarial Networks (GANs) are typically described as a two-player game between a generator and a discriminator, where the generator aims to produce realistic data, and the discriminator tries to distinguish between real and generated data. However, this setup often leads to mode collapse, where the generator produces limited variations in the data, failing to capture the full range of the target data distribution. This paper proposes a new training system to mitigate the mode collapse problem. Specifically, it extends the traditional two-player game of GANs into a multi-player game and introduces a peer-evaluation method to effectively train multiple GANs. In the peer-evaluation process, the generated samples from each GANs are evaluated by the other players. This provides external feedback, serving as an additional standard that helps GANs recognize mode failure. This cooperative yet competitive training method encourages the generators to explore and capture a broader range of the data distribution, mitigating mode collapse problem. This paper explains the detailed algorithm for peer-evaluation based multi-GANs training and validates the performance through experiments.

Keywords : Generative Models, Generative Adversarial Networks, Mode Collapse, Peer-Evaluation, Multi-Model Training

모드 붕괴를 완화하기 위한 다중 GANs 훈련 시스템

심 주 용[†] · 최 진 성^{††} · 김 종 국^{†††}

요 약

생성형 적대 신경망(GANs)은 보통 생성자와 판별자 사이의 두 플레이어 게임으로 설명된다. 여기서 생성자는 실제에 가까운 데이터를 생성하는 것을 목표로 하고, 판별자는 실제 데이터와 생성된 데이터를 구별하려고 한다. 하지만 이 방식은 종종 생성자가 데이터를 제한적으로 생성하여 데이터 분포의 다양성을 제대로 포착하지 못하는 모드 붕괴(mode collapse)로 이어질 수 있다. 이 논문에서는 이러한 모드 붕괴 문제를 완화하기 위한 새로운 훈련 시스템을 제안한다. 구체적으로, 기존의 이중 플레이어 게임을 다중 플레이어 게임으로 확장하고, 여러 GANs를 효과적으로 훈련시키기 위해 동료 평가(peer-evaluation) 방법을 제안한다. 동료 평가 과정에서는 각 GAN이 생성한 샘플들을 다른 플레이어들이 평가한다. 이는 외부 피드백을 제공하여 GAN이 모드 붕괴를 인식할 수 있는 추가적인 기준이 된다. 이러한 동료 평가 방법을 적용한 협력적이면서도 경쟁적인 다중 플레이어 게임 방식의 훈련은 생성자들이 데이터 분포의 더 넓은 범위를 탐색하고 포착하도록 돕는다. 이 논문에서는 여러 GANs를 효과적으로 훈련시키기 위한 알고리즘을 자세히 소개하고, 실험을 통해 그 성능을 검증한다.

키워드 : 생성 모델, 생성형 적대 신경망, 모드 붕괴, 동료 평가, 다중 모델 훈련

1. Introduction

The explosive popularity in generative tasks and applications has led to remarkable progress in generative model research[1-3]. Despite these developments, generating high-quality samples at fast-speed without falling into mode collapse problem at once is still challenging issue[4]. Among the different types of generative models,

this research specifically focuses on Generative Adversarial Networks (GANs)[5], which are particularly suffering from mode collapse problem, where the model generates only a limited subset of data or repeatedly produces the same samples[6].

GANs are usually defined as a two-player min-max game of two neural networks: a discriminator and a generator. These two are trained in an adversarial way, where the objective function is defined as:

$$\min \max V(D, G) = \mathbb{E}_{x \sim p_{data}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log (1 - D(G(z)))] \quad (1)$$

[†] 준 회원 : 고려대학교 정보통신기술연구소 박사후연구원

^{††} 비 회원 : 고려대학교 전기전자공학과 석·박통합과정

^{†††} 종신회원 : 고려대학교 전기전자공학과 교수

Manuscript Received : August 30, 2024

Accepted : September 22, 2024

* Corresponding Author : Jong-Kook Kim(jongkook@korea.ac.kr)

where $x \sim p_{data}$ is an image sample from the real data distribution, and $z \sim p_z$ is a prior input noise, which is generated randomly. Here, the generator G tries to minimize the objective function, generating realistic samples which are indistinguishable from the discriminator D . Whereas the discriminator D tries to maximize the objective function, identifying the fake samples generated by the generator G from the real data. Ideally, it reaches the equilibrium point at the global optimal state, where $p_G \sim p_{data}$. In this state, the generator G generates perfect samples and the discriminator D is unable to distinguish the real and generated distribution, i.e., $D(x) = \frac{1}{2}$.

The mode collapse problem usually occurs when the generator G captures only a few modes of data distribution that can easily deceive the discriminator D . In this state, the neither the generator G nor the discriminator D can make further improvements, as the discriminator fails to effectively differentiate between real and generated data. Resulting the objective function reaches a point where the loss does not change, indicating that the model is stuck in a suboptimal state without meaningful progress.

This work presents the multi-player GANs training process to mitigate mode collapse problem. The proposed method simultaneously trains multiple GANs by using an algorithm that includes a peer-evaluation process and an auxiliary training update. In the peer-evaluation process, the samples generated by each GANs are evaluated by their peers, which are the other GANs in the system. Following the peer-evaluation, an auxiliary training update is performed, using the values obtained from the peer-evaluation. This provides a valuable guidance based on external standards during the training process. The peer-evaluation helps in identifying the best-performing GANs. The score is calculated based on their peer-evaluation and the highest-scored GANs is selected as the best GANs. Each GANs then refers to the best GANs when updating its parameters during auxiliary training. The assumption is that the discriminator associated with the best-performing GAN is more accurate than the others, providing a reliable reference value during the training phase. Note that a preliminary version of the proposed method in this paper was published in our previous work[7, 8].

2. Related Works

There have been numerous efforts to solve the mode collapse problem of GANs. Some of the approaches focus on improving the learning process of GANs[9-12] and some works enforce GANs to cover diverse modes[13, 14]. Unrolled GANs[9], for example, reduce mode collapse by predicting how the discriminator will be affected in the next k steps to the current update of the generator. This process allows generator to learn a broader range of outputs well across the potential future states of discriminator, which reduce the likelihood of overfitting to a particular discriminator and helps the generator to cover more modes of data distribution. WGANs[10, 11] and LSGANs[12] use different divergence metrics to stabilize the training of GANs and generate better samples. InfoGAN[13] tried to allay the diversity issue by enforcing GANs to cover diverse modes. Unlike the classical GANs model which uses single input latent variable z , it adds latent code c to the input as (z, c) . Then, it extends information theoretic regulation, which tries to increase the Mutual Information (MI) between the code c added in the latent space and the generated sample. This helps to learn disentangled presentations, preventing mode collapse. ModeGAN[14], another trial to enforce diverse modes of GANs, trains generators in collaboration using encoders. It supports sample diversity by using the fact that the fake sample generated by the generator is likely to belong to the same mode when the real sample passes via the encoder-decoder. This work belongs to both approaches, in which it tries to improve the learning process of GANs by introducing an additional auxiliary training process and enforcing diverse modes by referencing multiple generators and discriminators.

This work is not the first attempt to extend the two-player game to the multi-player game[15]. There were previous researches that used multiple generators or discriminators[16-20]. MAD-GAN[16] uses multiple generators that capture different modes and the discriminator is trained to identify whether the samples are real or fake and find which generator has generated the fake sample. D2GAN[18] uses two discriminators and one generator. The additional discriminator is trained by using reverse KL distance, learning to judge fake data as real. It allows parameters to adjust the quality and diversity of the gen-

erated images by reducing the KL distance and reverse KL distance by minimizing the loss function. MCL-GAN[19] also uses multiple discriminators where each sample is assigned to best-suited discriminators. This makes each discriminator to be the expert model for assigned samples and helps to mitigate the mode collapse problem. This work uses multiple GANs to improve mode coverage and diversity performance, but it is different in that the proposed algorithm can be easily applied without any special structural changes, and multiple models evaluate values of each other.

3. Methods

In this system, a set of N pairs of GANs, each consisting of a generator G and a discriminator D , is considered :

$$U = (G_1, D_1), (G_2, D_2), \dots (G_N, D_N) \quad (2)$$

Each GANs is trained to maximize performance using two-stage update algorithm. The first-step update is conducted independently for each model, a process termed "individual training". This update process is applied across all pairs of GANs following the standard GANs training. For instance, when using vanilla GANs, each GANs is updated based on the objective function described in (1). Various types of GANs can be incorporated by simply adopting their respective objective function or algorithms in this stage. The second-step update is "peer training", which utilizes external value obtained from peer-evaluation process. In the peer-evaluation, GANs evaluate each other and the select best-performing GANs at each training step. Then, GANs update their parameters using a relative loss, calculated by referencing the best GANs. This approach is designed to calibrate the performance of the GANs, ensuring that each GANs improves by adjusting its output based on the relative loss in relation to the top-performing models. The overall algorithm is described in Algorithm 1.

3.1 Peer-evaluation

In the peer-evaluation process, each generator produces a set of samples, and these are evaluated by the discriminators, except the one paired to the generator. This approach ensures that the evaluation is not biased by

the paired discriminator, which may have adapted too closely to its corresponding generator. As a result, this process provides a more accurate assessment for each generator by having multiple discriminators to evaluate the generated outputs. The overall concept is described in Fig 1. Peer-evaluation follows three steps :

Algorithm 1. Overall Algorithm

Require :

Set of N GANs $U = (G_1, D_1), (G_2, D_2), \dots (G_N, D_N)$

K : batch size

```

while not converged
e
  for all  $(G_i, D_i) \in U$  do individual training
    for  $i := 1$  to  $N$ 
      sample noise  $\{z_k\}_{k=1}^K \sim p_z$ 
       $x_k^i \leftarrow G_i(z_k)$  ►Generate samples
      for  $j := 1$  to  $N$ 
        if  $j \neq i$  then  $v_k^{ij} \leftarrow D_j(x_k^i)$ 
        ►Evaluate samples
         $V^{ij} = \frac{1}{K}(v_1^{ij}, v_2^{ij}, \dots, v_k^{ij})$ 
        ►Peer-evaluation
      for all  $(G_i, D_i) \in U$  calculate  $S(G_i, D_i)$ 
       $(G^*, D^*) \leftarrow (G_1, D_1)$ 
      if  $S(G^*, D^*) < S(G_i, D_i)$  then
         $(G^*, D^*) \leftarrow (G_i, D_i)$ 
      else pass ►Select Best-GANs
    for all  $(G_i, D_i) \in U$  do
      sample data  $\{x_k\}_{k=1}^K \sim p_{data}$ 
      sample noise  $\{z_k\}_{k=1}^K \sim p_z$ 
       $\bar{x}_k \leftarrow G_i(z_k)$ 
       $L_{aux} = MSE(L_D(X, \bar{X}), L_{D^*}(X, \bar{X}))$ 
      update using  $L_D = L_D + \lambda L_{aux}$ 
      ►Auxiliary Update

```

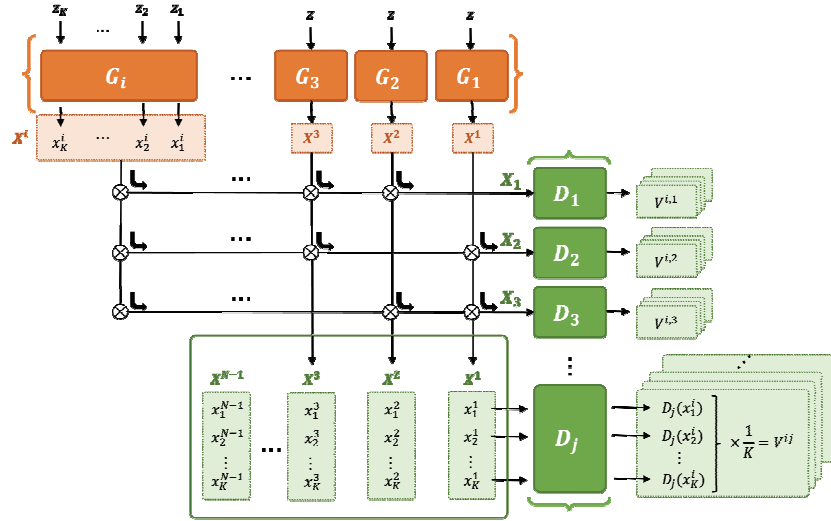


Fig. 1. Peer-Evaluation Concept

- i. (Sample Generation.) Each generator G_i generates a set of K samples based on its current parameters. The set of samples are denoted as $X^i = \{x_1^i, x_2^i, \dots, x_K^i\}$, where x_k^i is k -th sample produced by the generator G_i using a random noise z_k , i.e., $x_k^i \leftarrow G_i(z_k)$.
- ii. (Sample Evaluation.) The generated samples are then evaluated by the other discriminators. Each discriminator D_j ($j=1, \dots, N, j \neq i$) evaluates the K samples generated by G_i , effectively acting as external evaluators. The discriminator D_j provides an evaluated value v_k^{ij} to each samples x_k^i , which is the output of the discriminator $D_j(x_k^i)$ when x_k^i is given as input, i.e., $v_k^{ij} \leftarrow D_j(x_k^i)$.
- iii. (Calculate Peer-Evaluation Value) The final peer-evaluation value V^{ij} for generator G_i given by D_j is calculated as the mean evaluated values for all samples, $V^{ij} = \frac{1}{K}(v_1^{ij} + v_2^{ij} + \dots + v_k^{ij})$. Here, v_k^{ij} is the evaluation value given by discriminator D_j to the sample x_k^i at step 2.

3.2 Best GANs Selection

Based on the peer-evaluation results, the score of the GANs (G_i, D_i) is determined by using a score function. This score function calculates the mean difference between : (i) the mean values received from all other discriminators for samples generated by the given generator

G_i and (ii) the mean values submitted by the given discriminator D_i for samples generated by all other generators. This score is designed to reflect both how well a generated samples are evaluated by other discriminators and how critically its paired discriminator evaluates samples from other generators. The score function is formulated as follows :

$$S(G_i, D_i) = \frac{1}{N-1} \sum_{j=1, j \neq i}^N V^{ij} - \frac{1}{N-1} \sum_{j=1, j \neq i}^N V^{ji} \quad (3)$$

Here, $\sum_{j=1, j \neq i}^N V^{ij}$ is the sum of the values evaluated by all other discriminators D_j ($j=1, \dots, N, j \neq i$) for the samples generated by G_i . This represents the sum of the probability values that other discriminators predicted that the data generated by G_i are from the real. Samples generated by relatively poor GANs are more likely to be determined as fake, that is, the value v_k^{ij} is more likely to be low, thus resulting in lower V^{ij} values. $\sum_{j=1, j \neq i}^N V^{ji}$ is the sum of the values that discriminator D_i evaluated the samples generated by all other generators G_j ($j=1, \dots, N, j \neq i$). This is the sum of probability values that D_i determined as the real for the samples generated by other generators. If the GANs has relatively poor performance, it is highly likely that samples generated by other GANs $x_k^j \leftarrow G_j(z_k)$ will be judged as relatively good

samples, that is, $D_i(x_k^j)$ is more likely to be high. Therefore, the worse the GANs, the higher the value of V^{ji} will be measured. Taken together, the score is designed as (2), and the better the performance, the higher the score. For every GANs, score is calculated and the GANs having the highest score will be selected as the best GANs (G^* , D^*).

3.3 Auxiliary Loss

The auxiliary loss is calculated using the determined best GANs (G^* , D^*). The key idea is to match the loss output of a discriminator to the loss of the best discriminator D^* . Therefore, the auxiliary loss L_{aux} is defined as follows:

$$L_{aux} = MSE(L_D(X, \bar{X}), L_{D^*}(X, \bar{X})), \quad (4)$$

where L_D is the loss function of discriminator D . X is a batch of real samples, and \bar{X} is a batch of generated samples. Final loss function is augmented using the auxiliary loss:

$$L_D = L_D + \lambda L_{aux}, \quad (5)$$

where λ is a coefficient to be determined.

4. Experiments

4.1 Settings

To demonstrate the effectiveness of the proposed method, it is applied to two classical GAN models: vanilla GANs and WGAN, and then compared to models trained without the proposed method. This comparison shows the improvements achieved by the proposed approach in addressing the mode collapse issue. The models are trained and validated using a 2D-Gaussian dataset, which is a synthetic mixture of Gaussian distributions comprising eight modes. This simple dataset is chosen because mode collapse is more evident in such a simple dataset. For all experiments, eight GANs are trained simultaneously. All settings are kept the same across the experiments, except for the training algorithm, to ensure fairness in the comparison.

4.2 Qualitative Evaluation

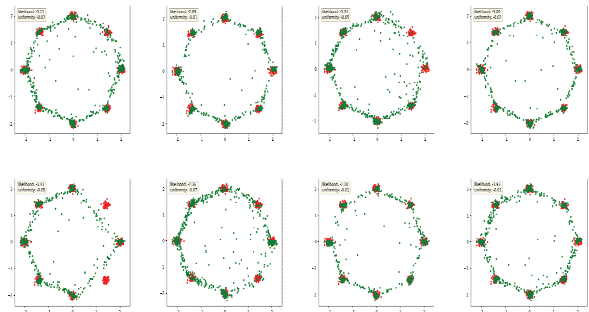


Fig. 2a. GANs Trained Without Using Proposed Method

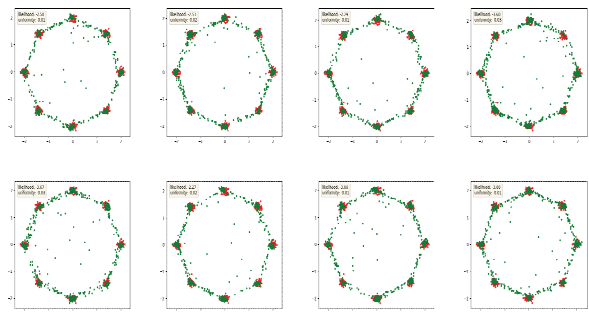


Fig. 2b. GANs Trained Using Proposed Method

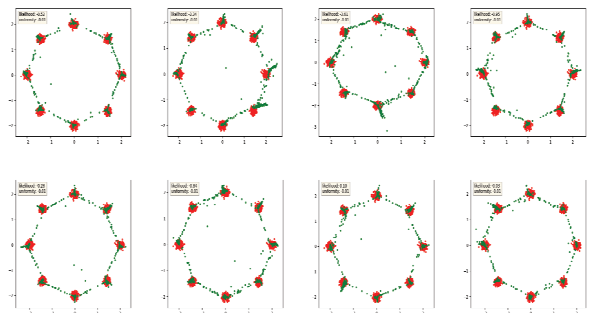


Fig. 3a. GANs Trained Without Using Proposed Method

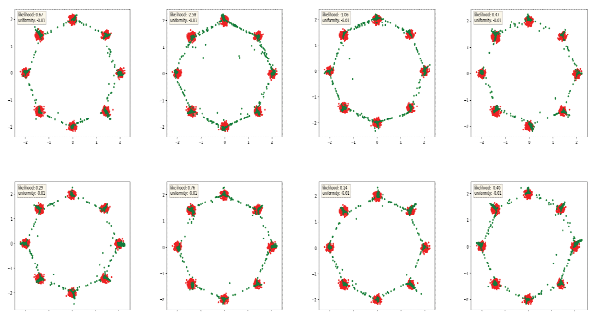


Fig. 3b. GANs Trained Using Proposed Method

In classical training, some GANs produced biased results, whereas the application of the proposed training method ensured that all modes were adequately covered.

The results are illustrated in Fig. 2. In Fig. 2a GANs trained using the standard approach often exhibit bias toward one side of the distribution or fail to capture data points in certain areas. In contrast, as depicted in Fig. 2b GANs trained using the proposed method successfully cover various points in the data distribution without experiencing mode collapse.

While the differences in performance are not as pronounced in the WGANs as they are in the vanilla GANs, the experiments on WGANs still demonstrate significant improvements. Fig. 3 shows the results using WGANs. When the proposed method was applied, the WGAN model effectively covered various modes of the data distribution in a balanced and sparse manner. Conversely, in WGANs trained without the proposed method, some samples tended to cluster on specific sides of the distribution, indicating a less balanced coverage.

4.3 Quantitative Evaluations

Quantitative metrics are implemented to assess both the quality and the mode coverage of the generated samples. For the performance evaluations, the likelihood metric was used. The larger value indicates better performance. There was a slight improvement. All values in Table 1 are the average values obtained from 10 experiments. The mean likelihood of eight GANs, and the worst performance results (min) among eight GANs are shown. The performance of the worst has improved particularly, which is expected to be due to the auxiliary training that refers to the best GANs, resulting in preventing it from going to failure mode.

In Fig. 4 and Fig. 5, the mode coverage is compared quantitatively. For GANs models, it can be seen that the proposed method is better at mode coverage. The higher the value, the higher the degree of spread of samples, which means there is less mode overfitting occurring. The differences between WGANs are not as clear as those of GANs, but there was slight progress in mode coverage.

Table 1. Performance Evaluation

Model	Likelihood	Classic	Proposed
GANs	mean	-3.530 ± 0.556	-3.235 ± 0.577
	min	-6.414 ± 1.528	-4.995 ± 1.120
WGANs	mean	-0.651 ± 0.831	-0.478 ± 0.571
	min	-4.756 ± 5.525	-2.691 ± 1.503




Fig. 4. Mode Coverage Results Compared Using GANs




Fig. 5. Mode Coverage Results Compared Using WGANs

5. Conclusion

In conclusion, this work introduces a novel approach to tackle the mode collapse problem in GANs by extending the traditional two-player game into a multi-player game. This framework effectively trains multiple GANs, offering a solution to the mode collapse problem. One notable contribution of this work is the introduction of peer-evaluation approach, which assists in selecting and valuing the best-performing GANs. This approach guides the training process effectively and helps prevent failure modes. Additionally, an auxiliary training process is proposed, providing proper guidance for training multiple GANs.

While the proposed method has successfully mitigated the mode collapse and instability problem, it does require additional computational efforts. Future work will focus on improving performance and implementing the approach using other GANs models. Exploring better score functions or alternative evaluation mechanisms may prove beneficial. Further investigation into updating using relative loss is also needed. All these limitations are left for future research.

References

- [1] A. Ramesh et al., “Zero-shot text-to-image generation,” in *Proceedings of the Advances in Neural Information Processing Systems*, pp.8821-8831, 2021.
- [2] M. Tao, B.-K. Bao, H. Tang, and C. Xu, “Galip: Generative adversarial clips for text-to-image synthesis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [3] H. Alqahtani, M. Kavakli-Thorne, and G. Kumar, “Applications of generative adversarial networks (gans): An updated review,” *Archives of Computational Methods in Engineering*, Vol.28, pp.525-552, 2021.
- [4] Z. Xiao, K. Kreis, and A. Vahdat, “Tackling the generative learning trilemma with denoising diffusion gans,” in *Proceedings of the International Conference on Learning Representations*, 2022.
- [5] I. Goodfellow et al., “Generative adversarial nets,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.27, pp.2672-2680, 2014.
- [6] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, “Improved techniques for training gans,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.29, 2016.
- [7] J. Y. Shim, J. S. B. Choe, and J.-K. Kim, “A study on auction-inspired multi-gan training,” in *Annual Spring Conference of KIPS*, Vol.30, No.1, pp.527-529, 2023.
- [8] J. Y. Shim, “Enhancements and applications of gans: Cross-modal generation, captcha system and mode collapse problem,” Ph.D. Dissertation, Department of Electrical and Computer Engineering, Korea University, 2024.
- [9] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein, “Unrolled generative adversarial networks,” *arXiv preprint arXiv:1611.02163*, 2016.
- [10] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *Proceedings of the International Conference on Machine Learning*, pp.214-223, 2017.
- [11] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, “Improved training of wasserstein gans,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.30, pp.5767-5777, 2017.
- [12] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, “Least squares generative adversarial networks,” in *Proceedings of the IEEE International Conference on Computer Vision*, pp.2794-2802, 2017.
- [13] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, “Infogan: Interpretable representation learning by information maximizing generative adversarial nets,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.29, 2016.
- [14] T. Che, Y. Li, A. Jacob, Y. Bengio, and W. Li, “Mode regularized generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations*, 2017.
- [15] M. Mohebbi Moghaddam et al., “Games of gans: game-theoretical models for generative adversarial networks,” *Artificial Intelligence Review*, pp.1-37, 2023.
- [16] A. Ghosh, V. Kulharia, V. P. Namboodiri, P. H. Torr, and P. K. Dokania, “Multi-agent diverse generative adversarial networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.8513-8521, 2018.
- [17] Q. Hoang, T. D. Nguyen, T. Le, and D. Phung, “MGAN: Training generative adversarial nets with multiple generators,” in *International Conference on Learning Representations*, 2018.
- [18] T. Nguyen, T. Le, H. Vu, and D. Phung, “Dual discriminator generative adversarial nets,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.30, 2017.
- [19] J. Choi and B. Han, “Mcl-gan: Generative adversarial networks with multiple specialized discriminators,” in *Proceedings of the Advances in Neural Information Processing Systems*, Vol.35, pp.29 597-29 609, 2022.
- [20] I. Albuquerque, J. Monteiro, T. Doan, B. Considine, T. Falk, and I. Mitliagkas, “Multi-objective training of generative adversarial networks with multiple discriminators,” in *Proceedings of the Advances in Neural Information Processing Systems, ser. Proceedings of Machine Learning Research*, K. Chaudhuri and R. Salakhutdinov, Eds., Vol. 97. PMLR, 09-15 Jun. pp.202-211, 2019.



심 주 용

<https://orcid.org/0000-0003-2828-9232>

e-mail : shimjoo@korea.ac.kr

Joo Yong Shim has been a postdoctoral scholar at the Department of Electrical and Computer Engineering, Korea University, Seoul, Republic of Korea, since March 2024, where she received her Ph.D. in electrical and computer engineering, in February 2024. She also received her B.S. in electrical engineering from Korea University, Seoul, Republic of Korea, in February 2019. Her current research interests include image generation, cross-modal generative models, and their applications to mobility systems.



최진성

<https://orcid.org/0000-0002-4865-5116>

e-mail : garangg@korea.ac.kr

Jean Seong Bjorn Choe has been a Ph.D. student at the School of Electrical Engineering at Korea University, Seoul, Republic of Korea, since 2018. He earned a B.S. degree in Atmospheric Science from the School of Earth and Environmental Sciences at Seoul National University in 2017. His research interests include reinforcement learning, game theory and artificial intelligence in video games.



김종국

<https://orcid.org/0000-0003-1828-7807>

e-mail : jongkook@korea.ac.kr

Jong-Kook Kim received the B.S. degree in electronic engineering from Korea University, Seoul, South Korea, in 1998, and the M.S. and Ph.D. degrees from the School of Electrical and Computer Engineering, Purdue University, in 2000 and 2004, respectively. He is currently a Professor with the School of Electrical Engineering, Korea University, where he joined, in 2007. He was with the Samsung SDS's IT Research and Development Center, from 2005 to 2007. His research interests include heterogeneous distributed computing, energy-aware computing, resource management, evolutionary heuristics, distributed mobile computing, artificial neural networks, deep learning and distributed robot systems. He is a Senior Member of the ACM.