

GENERATIVE AI FOR DEEP REINFORCEMENT LEARNING: FRAMEWORK, ANALYSIS, AND USE CASES

Geng Sun, Wenwen Xie, Dusit Niyato, Fang Mei, Jiawen Kang, Hongyang Du, and Shiwen Mao

ABSTRACT

As a form of artificial intelligence (AI) technology based on interactive learning, deep reinforcement learning (DRL) has been widely applied across various fields and has achieved remarkable accomplishments. However, DRL faces certain limitations, including low sample efficiency and poor generalization. Therefore, in this article, we show how to leverage generative AI (GAI) to address these issues and enhance the performance of DRL algorithms. We first introduce several classic GAI and DRL algorithms and demonstrate the applications of GAI-enhanced DRL algorithms. Then, we discuss how to use GAI to improve DRL algorithms from the data and policy perspectives. Subsequently, we introduce a framework that demonstrates an actual and novel integration of GAI with DRL, that is, GAI-enhanced DRL. Additionally, we provide a case study of the framework for UAV-assisted integrated near-field/far-field communication to validate the performance of the proposed framework. Moreover, we present several future directions. Finally, the related code is available at: <https://xiewenwen22.github.io/GAI-enhanced-DRL>.

INTRODUCTION

Deep Reinforcement Learning (DRL) is a transformative approach in the field of Artificial Intelligence (AI) that offers promising solutions to complex decision-making problems in various domains. DRL incorporates the principles of deep learning and reinforcement learning, where the introduction of neural networks has made DRL be capable of efficiently dealing with high-dimensional state and action spaces. Additionally, a DRL agent is able to learn directly from raw inputs without the need for artificial features and iteratively refine their strategies through trial-and-error interactions with an environment. This breakthrough capability has led to the extensive applications of DRL across diverse domains such as controlling robotic systems in dynamic environments.

Despite its remarkable achievements, traditional DRL algorithms suffer from certain limitations. One of the main challenges is the low sample efficiency for training deep neural networks in reinforcement learning environments. As the process

of learning optimal policies often requires long and extensive interaction with an environment, data collection for DRL may be costly or impractical in the real world. Moreover, DRL models typically exhibit limited generalization capabilities, making it difficult to extend learned policies to unseen environments or tasks.

Fortunately, Generative AI (GAI) can offer a promising solution to the challenges faced by DRL. Unlike discriminative AI, which aims to learn boundaries between different data categories, GAI can learn the distribution of data to capture its latent structure. Therefore, GAI is widely used in content analysis and creation domains. For example, ChatGPT (<https://openai.com/index/chatgpt>) has revolutionized human-computer interaction by leveraging large language models (LLMs). Additionally, the recently introduced video generation model Sora (<https://openai.com/index/sora>) has sparked a trend in the field of video generation. The widespread adoption of GAI can be attributed to several key advantages over other AI methods.

High-quality Data Generation Ability: GAI can generate new credible data samples by learning the distribution of the data, which helps to solve the problem of data scarcity.

Feature Extraction Ability: GAI can extract useful features or representations from raw data by learning the distribution and latent structure of the data, which helps to improve the data analysis ability.

Generalization Ability: GAI can effectively apply the knowledge and representations learned from one task to other tasks, which accelerates the learning process and improves the performance in new environments.

Benefiting from the aforementioned advantages, GAI can be utilized to enhance DRL algorithms due to higher sample efficiency, stronger generalization capabilities, and robustness to environmental changes. For instance, a DRL agent may struggle in some specific environments, where the agent is not allowed to engage in extensive trial-and-error interactions within the environment. In this case, GAI can learn from limited training data to generate synthetic data for DRL training, which contributes to alleviating data scarcity and accelerating the learning process. Additionally,

Geng Sun is with Jilin University, China, and also with Nanyang Technological University, Singapore; Wenwen Xie and Fang Mei (corresponding author) are with Jilin University, China; Dusit Niyato is with Nanyang Technological University, Singapore; Jiawen Kang is with Guangdong University of Technology, China; Hongyang Du is with The University of Hong Kong, China; Shiwen Mao is with Auburn University, USA.

Digital Object Identifier: 10.1109/MWC.001.2400176

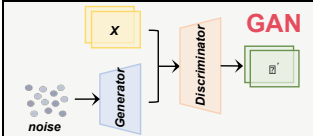
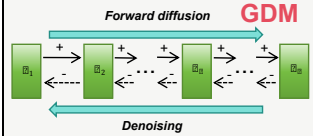
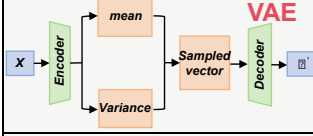
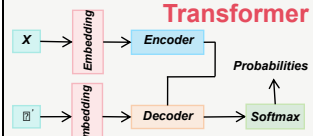
Model	Components	Pros and Cons	Applications	Capabilities to improve DRL
 <p>GAN</p>	<ul style="list-style-type: none"> ◆ Generator : Generate fake data that is similar to real data ◆ Discriminator : Judge the authenticity of data 	<ul style="list-style-type: none"> ✓ Generate high-quality and realistic data samples ✓ Don't require a large amount of annotated data ✗ Mode collapse 	<ul style="list-style-type: none"> ➢ Video generation ➢ Image restoration ➢ Virtual reality ➢ Super resolution 	<p>Generative Capability: GAI can facilitate the generation of diverse data samples for DRL, which can effectively mitigate the challenge of data scarcity, particularly in certain specific scenarios where obtaining sufficient real-world data is difficult or expensive.</p>
 <p>GDM</p>	<ul style="list-style-type: none"> ◆ Forward : Gradually inject noise into input ◆ Denoising : Remove noise to approach real distribution 	<ul style="list-style-type: none"> ✓ Easy to understand and implement ✓ Capture true distribution ✓ Enhanced denoising ✗ Long inference time 	<ul style="list-style-type: none"> ➢ Image denoising ➢ Image generation ➢ Data augmentation ➢ Provide high quality solutions for problem 	<p>Exploratory Capability: GAI can enhance the policy network of DRL, which can help the DRL agent explore the environment more comprehensively and make more informed and effective decisions.</p>
 <p>VAE</p>	<ul style="list-style-type: none"> ◆ Encoder : Encode the input data into a latent space representation ◆ Decoder : Decode the latent space 	<ul style="list-style-type: none"> ✓ Learning the latent space of data ✓ Suitable for denoising and reconstruction ✗ Training instability 	<ul style="list-style-type: none"> ➢ Speech synthesis ➢ Image reconstruction ➢ Unsupervised learning ➢ Feature learning ➢ Dimension reduction 	<p>Extraction Capability: GAI can extract useful features to provide DRL with richer and more informative input data, which helps to reduce the learning difficulty of DRL in high-dimensional scenarios. As such, This process can enhance the efficiency of the learning process.</p>
 <p>Transformer</p>	<ul style="list-style-type: none"> ◆ Encoder : Encode the input to capture semantic information ◆ Decoder : Convert the hidden representation to the target sequence 	<ul style="list-style-type: none"> ✓ Parallel processing of input sequences ✓ Highly interpretable ✓ Efficient multi-modal interaction ✗ High computational cost 	<ul style="list-style-type: none"> ➢ Machine translation ➢ Text generation ➢ Language modeling ➢ Speech recognition ➢ Sentiment analysis 	<p>Adaptability: GAI can learn shared features and patterns across different tasks and transfer this knowledge to new tasks in DRL, which can improve the learning efficiency.</p>

FIGURE 1. Summary of typical GAI models.

GAI can provide potential representations of data, which helps a DRL agent explore the environment more efficiently and generalize a policy better to unknown states. Therefore, the introduction of GAI brings new opportunities to DRL.

The combination of GAI and DRL is not straightforward and faces certain key issues. Specifically, GAI contains a variety of models that adept in solving different problems. Therefore, how to select an appropriate GAI model and integrate it with a specific DRL algorithm to solve certain problems is important and challenging. This is because the process involves the need for extensive studies and experimental validation of the compatibility between different models and algorithms to unleash full potential of GAI and DRL.

Motivated by this, we attempt to propose a comprehensive tutorial. *To the best of our knowledge, this is the first work to comprehensively summarize how DRL algorithms can be enhanced using various GAI models.* The contributions can be summarized as follows:

- We present the principles of important and widely used GAI models and DRL algorithms. Then, we briefly describe several applications of DRL enhanced with GAI.
- We explore different integration of GAI in DRL, focusing on how DRL can be enhanced from a data and strategy perspective. Moreover, we comprehensively summarize the advantages and disadvantages of the GAI-enhanced DRL algorithm.
- We propose a novel framework which deeply demonstrates the integration process of several common GAI models individually in the DRL algorithm. Moreover, we show the role of integrating multiple GAI models simultaneously in DRL to solve different challenges in DRL algorithms. Furthermore, we construct a case study to demonstrate the effectiveness of the proposed framework.

OVERVIEW OF GAI AND DRL

In this section, we introduce the basics of GAI and DRL.

GAI MODELS

The rapid development of GAI relies on several promising techniques, and we focus on the most commonly used Generative Adversarial Network (GAN), Generative Diffusion Model (GDM), Variational Autoencode (VAE), and Transformer. We summarize the technical details of GAI models and their comparison in Fig 1.

GAN: GAN is a deep learning model composed of a generator and a discriminator. The former aims to generate samples similar to real data, while the latter attempts to distinguish between samples generated by the generator and real data. By adversarially training these two networks, the credibility of generated samples is gradually enhanced. The applications of GAN include image generation and style transfer.

GDM: The training of GDM consists of forward process and reverse process. Specifically, GDM deconstructs training data by gradually adding Gaussian noise in the forward process, then reverses this process to generate the desired data from the noise, which simulates the evolution of the data. The applications of GDM include image restoration and solving optimization problems.

VAE: VAE consists of an encoder and a decoder. The former maps input data to a probability distribution in the latent space, while the latter reconstructs samples from the latent space into the original data. Clearly, the higher the similarity between the original and reconstructed samples, the higher the model accuracy. VAE is widely applied in fields such as image generation and feature learning.

Transformer: Transformer is a deep learning model based on the self-attention mechanism, consisting of an encoder and a decoder. Specifically, the

DRL Algorithm	Analysis		
	Policy	Key Features	Pros & Cons
DQN	Off-policy	Estimating Q-function using DNN and handle discrete action space	<ul style="list-style-type: none"> ✓ Simple to implement ✓ Good generality × Training instability
DDPG	Off-policy	Deterministic policy gradient method	<ul style="list-style-type: none"> ✓ Good stability × Poor convergence
TD3	Off-policy	Dual-Q network and delayed policy network updates	<ul style="list-style-type: none"> ✓ Improved training stability ✓ Reduced overfitting × Hyperparameter sensitivity
PPO	On-policy	Limiting the magnitude of policy network updates	<ul style="list-style-type: none"> ✓ High stability ✓ High training efficiency × Low sample efficiency
SAC	Off-policy	Entropy regularization terms and dual-policy networks	<ul style="list-style-type: none"> ✓ Explore action space more fully ✓ Insensitive to parameters × High training complexity

TABLE 1. The analysis of different DRL algorithms.

principle of the Transformer is to capture the dependencies between different positions in a sequence through self-attention mechanism, while enhancing representational ability of the model through multi-head attention mechanism. Transformer is widely used in the field of natural language processing, such as machine translation, language modeling.

DRL ALGORITHMS

Reinforcement Learning (RL) is a machine learning method that aims to enable an agent to maximize cumulative rewards through trial-and-error learning during their interaction with the environment. Specifically, the agent learns to make optimal decisions in a given environment based on the state and feedback rewards of the environment. Note that traditional RL algorithms use tables to estimate the Q-function that helps the agent to improve policy. However, it is clearly not suitable for high-dimensional and large state and action spaces. As such, DRL combining deep neural networks (DNN) and RL has been proposed, where DNN is used as an approximator of the Q-function. We summarize several representative DRL algorithms, namely, deep Q-network (DQN), deep deterministic policy gradient (DDPG), twin delayed deep deterministic policy gradient (TD3), proximal policy optimization (PPO), and soft actor-critic (SAC). The details of the DRL algorithms above are shown in Table 1.

Considering the ability of DRL to handle dynamic problems and find optimal solutions in the absence of information, it has achieved remarkable results in data communications and networking. However, DRL still has some limitations that may be mitigated by GA.

Sample Inefficiency: DRL typically requires an extensive interaction with the environment to learn effective policy. In this case, GAI (e.g., GAN) can help generate synthetic data to augment training dataset, which improves sample efficiency [1].

Challenges in Complex Environment: A DRL agent may fall into local optima or struggle to explore large and complex environments efficiently. In this instance, GAI (e.g., VAE) enables the agent to analyze the underlying data distribution

and perform feature extraction, which helps to accurately model the environment and improve the stability of training process [2].

Instability and Low Learning Efficiency: DRL may suffer from instability and balance of exploration-exploitation trade-off. In this situation, GAI (e.g., GDM) can be used to improve the policy network of DRL, which helps to enhance the stability of exploration [3].

APPLICATIONS OF GAI-ENHANCED DRL

Thanks to the excellent generative and analytical capabilities of GAI, it has been widely applied to enhance DRL in various fields. We have summarized typical applications.

Signal Processing: Training an agent to solve control tasks directly from multimodal signals has been shown to be difficult. In [4], a framework combining Transformer and DRL is proposed, where Transformer is used for feature extraction of multimodal data, which can better integrate important information carried by different signals to assist the agent in making more accurate decisions.

Wireless Communications: GAI can be utilized to ameliorate the problem of high training cost and low sample efficiency against DRL in wireless environments. In [5], the authors proposed a GAN-based auxiliary training mechanism that reduces the overhead of interacting with the real-world by generating environment states for offline training.

Edge-Based AIGC Service: GAI combined with DRL is utilized to provide edge-based AIGC services. For example, the authors in [3] proposed a diffusion model based approach to generate optimal AIGC service provider selection decisions and combined it with SAC algorithm to improve the efficiency and effectiveness of the algorithm.

Sensor Networks: GAI can be utilized to improve the accuracy of DRL decision making for cognitive internet of things network. In [6], GAI is employed to enhance the value network of DRL by improving the stability and accuracy of action value estimation through adversarial training, thereby enhancing the decision-making performance of DRL.

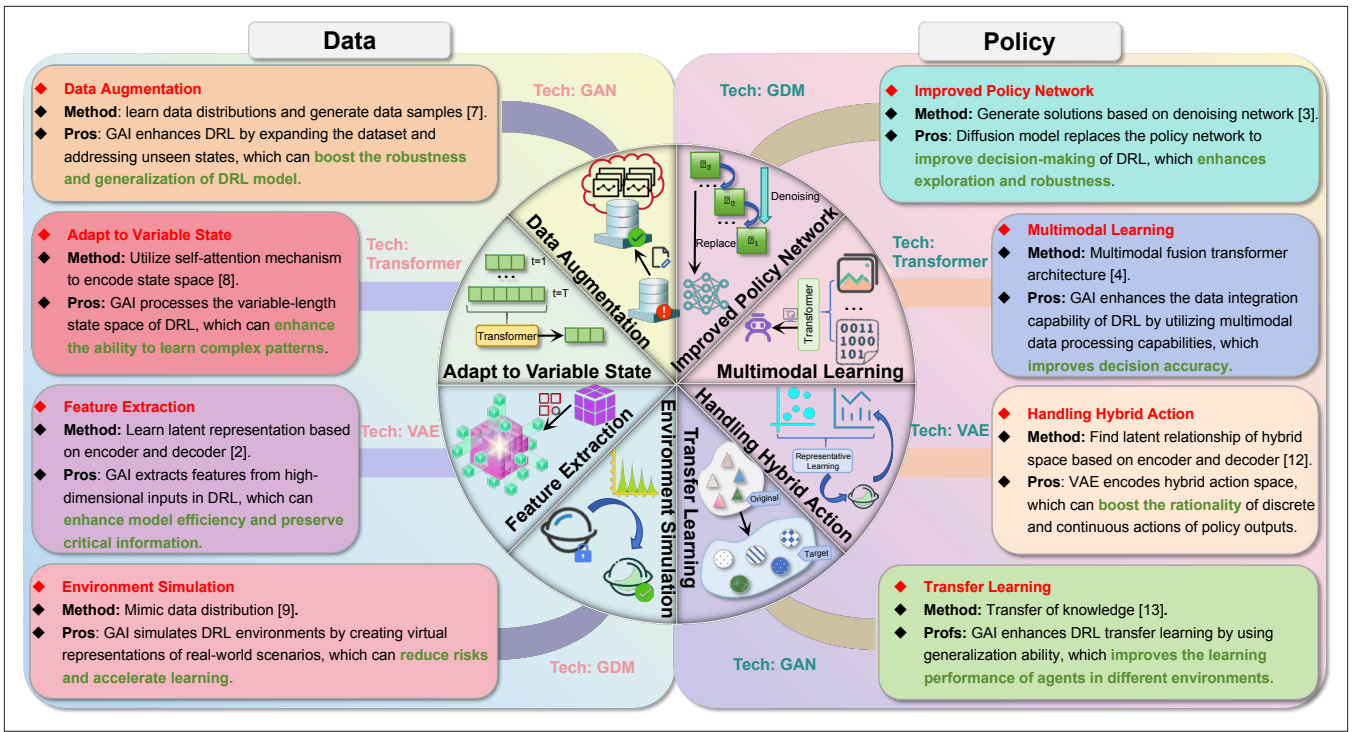


FIGURE 2. Summary of data and policy performance improved by GAI in DRL.

Lesson Learned: From the above analysis of the algorithms and applications, we observe that GAI models can effectively address the limitations of DRL. Specifically, the data generation capability of GAN contributes to tackling the data scarcity issue in certain scenarios. VAE can be employed for extracting features of high-dimensional data. Moreover, the exploration capability of GDM can be utilized to enhance the policy network of DRL, thereby improving the quality and stability of decision-making. The proficiency of Transformers in handling multimodal and variable-length data expands the scope of applications for DRL.

GAI-ENHANCED DRL

In this section, we introduce how GAI can improve DRL from the data and policy perspective. The summary of improvements are shown in Fig. 2.

FROM DATA PERSPECTIVE

Based on aforementioned unique features, GAI can effectively address the data-related challenges to improve the convergence speed and solution accuracy of DRL.

Data Augmentation: In most DRL training processes, agents sample experiences stored in an experience replay buffer randomly or in a specific way (e.g., prioritized experience replay) and utilize these experiences to update neural networks for better performance. Simultaneously, it is challenging for agents to output high-quality actions based on unobserved states or with limited data in such training approach. Adopting GAI to augment data can be a solution. For example, in [7], the authors proposed a GAN-powered multi agent DDGP (MADDPG) framework for joint optimizing terminal-cooperative caching and offloading for virtual reality tasks. Specifically, given high costs and privacy risk issues associated with collecting data from multiple users in a decentralized network, GANs

are used to learn the distribution of real data and generate appropriate data for each agent based on the state. Then, each agent inputs the corresponding state samples generated by GAN into the actor-critic network and stores the resulting experiences in the experience buffer for sharing, which augments the agent with limited samples of real virtual reality data. Numerical results indicate a 5.32% enhancement in energy efficiency for GAN-MADDPG compared with MADDPG trained on a real environment, with similar performance in terms of task completion latency.

Adaptation to Variable State: DRL with a standard deep neural network typically deals with fixed-length state spaces, which makes it challenging to handle problems with variable-length state spaces. However, in certain scenarios such as edge computing and task offloading, dynamic state spaces are relatively common, which limits the applicability of DRL in such contexts. Note that GAI is excellent in dealing with variable-length inputs, especially Transformer. A Transformer-based multi-agent DRL framework for scalable multi-unmanned aerial vehicle (UAV) area coverage was proposed [8]. In such a scenario, a deep neural network based on a fixed-dimension multilayer perceptron (MLP) is no longer applicable. This is because when the number of UAVs varies, some dimensions of the state space are blank, which may lead to neural network errors. In this case, Transformer is adopted to encode the state space and assign different weights to deal with variable input dimensions through a self-attention mechanism, which can mitigate the effect of blank states on decision making. Meanwhile, there is a large amount of heterogeneous information in multi-UAV scenarios, where the importance of this information varies. The self-attention mechanism can dynamically learn the importance of

By fully leveraging its ability to extract latent patterns and adapt to various data types, GAI can enhance the decision-making capability of DRL, enabling it to tackle complex tasks more efficiently and flexibly.

different positions in the input sequence, which empowers the agent to extract most relevant information. Experimental results show that Transformer-based DDPG achieves the best network performance. Compared with the expert coverage-first algorithm, the Transformer-based DDPG algorithm improves the average coverage score by 39.7% and the fairness index by 45.3%.

Feature Extraction: When dealing with high-dimensional data, DRL requires a significant amount of resources and time. Therefore, it is important for DRL to reduce the dimension of data without losing important information. Fortunately, GAI can analyze the latent relationships in data and perform feature extraction. Therefore, GAI can be used to process input data of DRL to alleviate the limitations of DRL when facing high-dimensional data. For instance, in [2], the authors proposed an environment perception framework for autonomous driving, where VAE is combined with SAC to focus on uninterrupted and reasonably safe autonomous driving. Specifically, VAE is used to process input driving environment images for feature extraction, which improves sample efficiency and facilitates learning processes with fewer samples but higher robustness. Simulation results on the DonKey simulator demonstrate the VAE-SAC method enables an autonomous vehicle to remain on track for the maximum time in a given timeframe. Note that although both VAE and Transformer can be used to handle the state space of an agent, their focuses are different. VAE [2] primarily aims at improving sample efficiency by extracting features to reduce the dimensionality of inputs. Conversely, Transformer [8] is not used explicitly for dimensionality reduction, and its main purpose is to handle variable-length state spaces through focusing on important parts of the input sequence.

Environment Simulation: Agent-environment interaction is a crucial aspect of DRL training processes. However, in certain scenarios, an agent are not allowed to perform trial-and-error training in real environments. Therefore, considering that GAI is able to learn and model data distribution, it can be harnessed to simulate the environment of DRL. For example, the authors in [9] proposed a primary-user-friendly dynamic spectrum anti-jamming framework. Note that it is inevitable for a second user (SU) to interfere the primary user (PU) even DRL is used by the SU to optimize spectrum access, which is undesirable in overlay cognitive radio network. Therefore, a GAN-based virtual environment (VE) is utilized to accurately simulate the spectrum environment. Subsequently, a channel decision network (CDN) based on DRL learns an optimal spectrum access strategy in VE through offline training. Experimental results show that the proposed framework converges much faster than the scheme that trains CDNs in the spectrum environment from scratch.

Lesson Learned: From the studies above, we summarize the uses of GAI to improve DRL from the data perspective as follows.

Data Generation: Given the powerful ability to mimic data distributions and generate credible data samples, GAI can be used to expand training datasets and anticipate unknown situations [7, 9].

Data Processing: GAN can be used to process dynamic data and high-dimensional data with unique analyzing and learning capabilities [2, 8].

By fully leveraging its ability to extract latent patterns and adapt to various data types, GAI can enhance the decision-making capability of DRL, enabling it to tackle complex tasks more efficiently and flexibly.

Improved Policy Network: The structure of a DRL policy network can affect quality of the output action. Currently, GAI has been used to enhance DRL policy network, especially in diffusion model. In [3], the authors proposed a diffusion model-based SAC algorithm to generate an optimal AIGC service provider (ASP) selection decision for better and broader AIGC services in wireless networks. Specifically, the diffusion model is utilized to replace the policy network of SAC to generate action distributions based on observed states. It is worth noting that optimization problems in wireless networks usually are difficult to achieve optimal solutions. Thus, only the reverse process of the diffusion model is employed to construct the policy network for SAC. Extensive experimental results demonstrate the effectiveness of the proposed diffusion model-based algorithm, which outperforms seven representative DRL algorithms (e.g., PPO) in both the ASP selection problem and various standard control tasks.

Multimodal Learning: DRL is typically limited to handling a single type of data, which restricts its applicability. The GAI's ability of processing multimodal data can be employed to enhance DRL. In [4], the authors proposed a Transformer-based DDPG to solve autonomous vehicle decision problem in complex scenarios. Multimodal Transformer is utilized to handle different types of data, that is, LiDAR point cloud and images. Then, DDPG is employed to complete the subsequent autonomous driving decision-making task based on the extracted features. The experimental results indicate that the DDPG based on multimodal Transformer achieves higher rewards than DDPG that uses only images, which suggests that multimodal Transformer plays a significant role in extracting key multi-modal features.

Handling Hybrid Action: DRL is typically used to deal with hybrid action spaces containing both discrete and continuous actions, especially for optimization problems involving scheduling. Current common practices either approximate the hybrid space by discretization or relax it into a continuous set. However, continuous actions and discrete actions typically interact with each other in general, while the aforementioned discretization or relaxation methods ignore this intrinsic relationship and leads to information loss. In this case, GAI constructs latent space representations and analyze the dependencies of hybrid action space. The authors in [10] proposed a novel VAE-based DRL framework called HyAR. The main idea of the proposed framework is to construct a unified and decodable representation space for the original discrete-continuous hybrid action, in which the agent learns a latent strategy. Then, the selected latent action is decoded back into the original hybrid action space to interact with the environment. The results show that HyAR has advantages over the baseline (e.g., PADDPG), especially in high-dimensional action space.

GAI DRL	GAI Models				Analysis
	Transformer	GDM	GAN	VAE	
DQN	<ul style="list-style-type: none"> Introduce the transformer structure into policy neural networks. 	<ul style="list-style-type: none"> Model in-support trajectory sequences. 	<ul style="list-style-type: none"> Apply the GAN structure to actor and critic networks [12]. 	<ul style="list-style-type: none"> Capture the low dimensional latent states for facilitating DQN agent learning. 	Potential Benefits: Expressiveness: GAI is able to learn complex sample distributions well, which enables it to be integrated into the policy network of DRL for decision-making.
DDPG	<ul style="list-style-type: none"> Process multimodal data [4]. Handle variable state space [8]. 	<ul style="list-style-type: none"> Introduce the GDM structure into policy networks [13]. 	<ul style="list-style-type: none"> Generate complete dataset including not appeared states [7]. 	<ul style="list-style-type: none"> Provide the extracted feature state as input for policy network [2]. 	
TD3	<ul style="list-style-type: none"> Introduce transformer module to the actor network. 	<ul style="list-style-type: none"> Approximate actions by noise networks to improve the speed of training. 	<ul style="list-style-type: none"> Generate environment states for offline training of TD3 agent [5]. 	<ul style="list-style-type: none"> Provide the latent representation of continuous-discrete hybrid action space [10]. 	Sample Quality: GAI has strong imitation and generation abilities, allowing it to generate highly credible new samples based on existing samples to expand the training dataset of DRL.
PPO	<ul style="list-style-type: none"> Employ transformer structure as the police network. 	<ul style="list-style-type: none"> Generate samples. 	<ul style="list-style-type: none"> Generate new training data based on historical data [1]. 	<ul style="list-style-type: none"> Encode input video to reduce the state dimension. 	
SAC	<ul style="list-style-type: none"> Extract latent features. Encode historical information and hybrid space. 	<ul style="list-style-type: none"> Apply GDM structure to policy network [3]. Simulate expert behavior to generate optimal policy. 	<ul style="list-style-type: none"> Provide a simulated environment for SAC agent [9]. 	<ul style="list-style-type: none"> Map high-dimensional states into the latent vectors. Encode tasks and infer tasks. 	Flexibility: The ability of GAI to model diverse behaviors is particularly useful in DRL, where the agent needs to adapt to a variety of situations and tasks.

TABLE 2. The use of GAI for various DRL algorithms.

Transfer Learning: One of the major dilemmas of DRL is that if the structure of network states change dramatically, we have to reformulate the problem and retrain the model from scratch. Fortunately, GAI can be used for knowledge transfer, which can alleviate the limitations. For example, in [11], the authors proposed a diffusion model-based transfer learning method to solve the limited data problem in image generation tasks. The proposed method trains a diffusion model with an adversarial noise selection and a similarity-guided strategy, which improves the efficiency of the transfer learning process. Extensive experiments in the context of few-shot image generation tasks demonstrate that the proposed transfer learning method is not only efficient but also excels in terms of image quality and diversity. Inspired by this work, we can adopt GAI model (e.g., diffusion model and GAN) to design transfer learning for DRL.

Lesson Learned: We find that the GAI approaches to improve policy performance can be summarized into two aspects.

Independent Strategy: GAI can be regarded as an independent strategy to preprocess the inputs of DRL, which can enhance the training efficiency of DRL [4].

Integrated Strategy: GAI can be integrated into DRL architecture to enhance the solution capability of DRL [3].

DISCUSSION

From the applications shown in Table 2, we can find that integrating GAI into DRL can significantly improve the performance of DRL, which can be summarized as follows.

Improved Data Diversity: The generative and learning capability of GAI can increase the diversity of training data for DRL to alleviate the data scarcity problem in certain particular scenarios, which

can improve the speed of DRL training. In addition, the multimodal data processing capability of GAI greatly broadens the application scope of DRL.

Enhanced Analysis Ability: Some GAI models, such as VAEs, are particularly effective for high-dimensional data, as it can eliminate the influence of redundant information and enhance the speed of DRL training.

Improved Decision Ability: The reverse process of the diffusion model is well-suited to the policy networks of DRL, where multi-round iterations and denoising processes help agents explore better actions. Additionally, GAI has inspired the implementation of transfer learning in DRL to facilitate interactions across similar domains.

Although incorporating GAI into DRL addresses some issues faced by DRL, it also introduces the following disadvantages.

High Computation Complexity: Integrating GAI models into DRL algorithms may introduce additional computational costs, such as tuning the network architecture of GAI models and a large number of hyperparameters, which can increase computational overhead during training. Moreover, the training process of GAI models requires iterative optimization, which leads to an increase in the overall training time of DRL.

High Resource Demand: GAI models typically require substantial memory to store model weights, parameters, and intermediate results, which is particularly pronounced in data-intensive applications. Moreover, GAI models often have complex structures and require high-performance computing resources to support their training when performing computationally intensive tasks.

Therefore, GAI-enhanced DRL is suitable for solving the issues of low sample efficiency and insufficient exploration by generating diverse training data and high-quality actions with sufficient

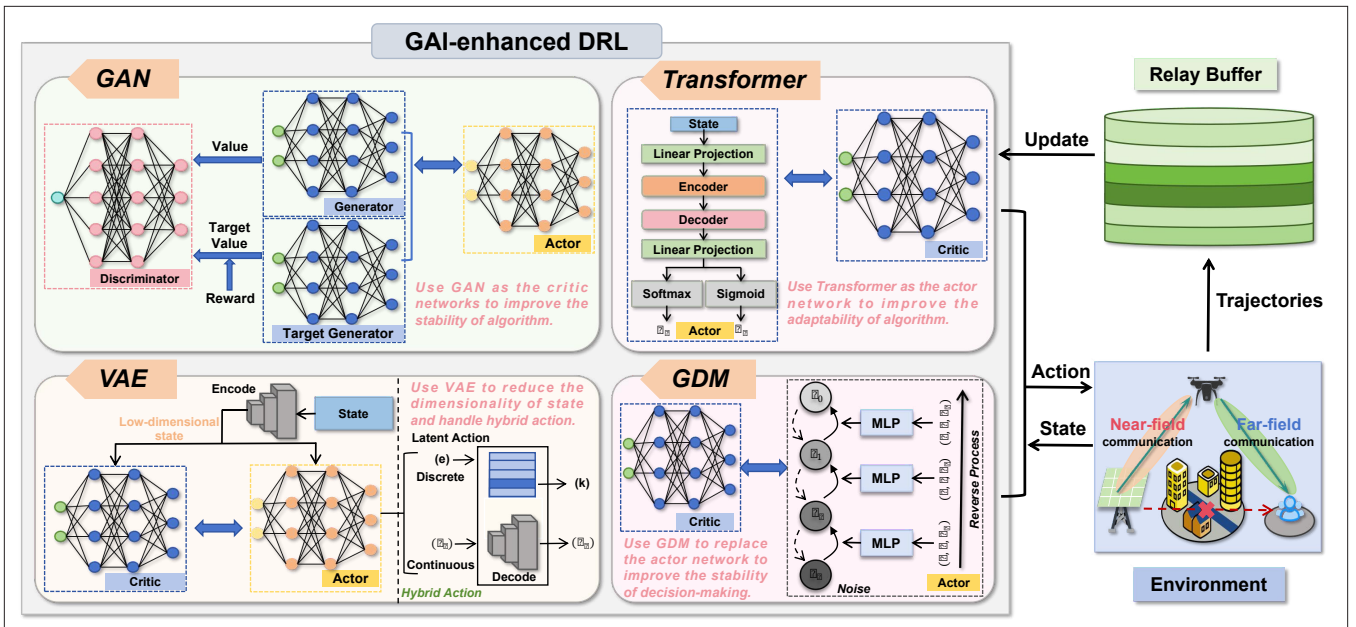


FIGURE 3. Proposed Framework of GAI-enhanced DRL on UAV-assisted integrated near-field/far-field communication.

computational resources.

CASE STUDY: GAI-ENHANCED DRL IN NEAR-FIELD COMMUNICATION

In this section, we first propose a novel framework for GAI-enhanced DRL. We then introduce a case study to demonstrate effectiveness of the proposed framework.

PROPOSED FRAMEWORK

As shown in Fig. 3, we investigate an integration of various GAI models into DRL algorithms to enhance the performance in different ways.

GAN-Enhanced DRL: We enhance the critic network of DRL by using GAN. Specifically, the generator network outputs estimated action values, while the target generator network obtains the target action values. The discriminator network attempts to minimize the distance between the estimated action values and the target action values calculated by the Bellman operator. Introducing adversarial learning mechanism inherent in GAN to approximate the action-value distribution avoids the negative impact of randomness on rewards, which enhances the stability of DRL.

VAE-Enhanced DRL: We use VAE to reduce the dimensionality of the high-dimensional state space in DRL. Specifically, we input the high-dimensional state data into the encoder of the VAE to extract a lower-dimensional latent variable. Subsequently, the decoder attempts to reconstruct the original state data from the latent variable. By minimizing the discrepancy between the reconstructed data and the original data, the VAE is able to maintain the quality of the lower-dimensional latent variable while performing dimensionality reduction. Additionally, VAE can construct a latent representation space for continuous parameters conditioned on state and embedding of discrete actions to handle hybrid actions.

Transformer-Enhanced DRL: We enhance the actor network of DRL by using Transformer. Specifically, we replace the MLP with a network

based on the attention mechanism of Transformer to analyze the current state in the environment. This enables Transformer-enhanced DRL to more effectively capture long-range dependencies within sequences, making it well-suited for decision-making tasks that require long-term memory or complex historical information.

GDM-Enhanced DRL: We improve the policy network of DRL by employing the reverse process of GDM. Specifically, we consider the policy network as a denoiser, progressively adding denoising noise to the initial Gaussian noise to recover or discover the optimal actions. Benefiting from the generative ability and expressiveness of GDM, it can generate diverse and high-quality actions, which makes it suitable for improving DRL algorithms with limited exploration capabilities.

SCENARIO DESCRIPTION

Large-scale antenna array has demonstrated an important enhancement in the spectrum efficiency of wireless systems. It is noteworthy that with the significant increase in carrier frequency and the number of antennas, the well-known Rayleigh distance expands to tens or even hundreds of meters. In this context, near-field communication becomes more prominent [14]. Unlike traditional far-field communication where the electromagnetic field is simply modeled by plane waves, the electromagnetic field in near-field communication needs to be accurately modeled by spherical waves. Therefore, we can exploit this unique channel characteristic to achieve flexible beamforming, concentrating the beam energy at specific spatial locations rather than focusing it in a specific direction as in traditional far-field communication. The characteristic above provides near-field communication with distinct advantages in improving signal quality, especially in densely populated urban areas.

In the case study, we considered a UAV-assisted ground communication, where the UAV serves as a relay to transmit data from a base station (BS) equipped with a large-scale antenna array to the

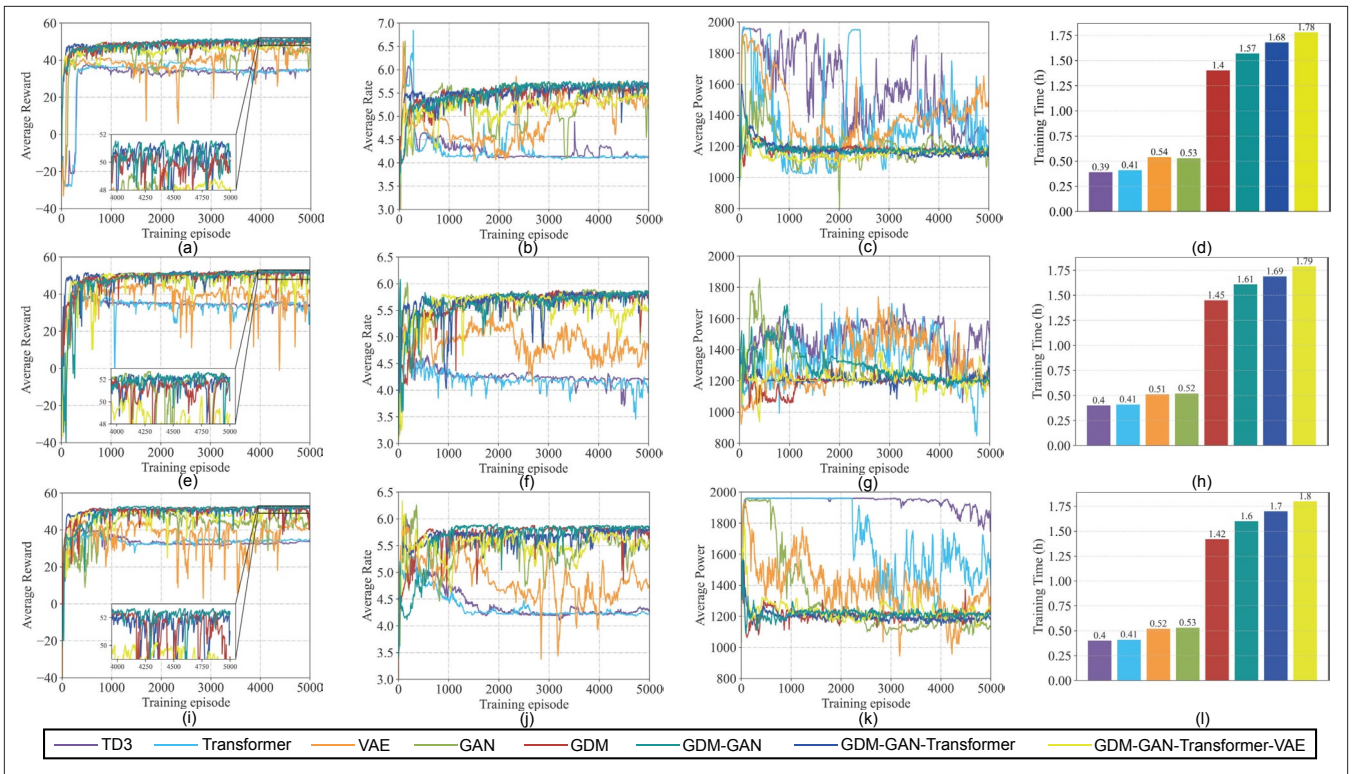


FIGURE 4. The convergence curves of rewards, rates, and power consumption, as well as the training time, of the case study under three different action spaces. Specifically, the results in continuous action space are shown in a), b), c), and d), respectively. The results in discrete action space are shown in e), f), g), and h), respectively. The results in hybrid action space are shown in i), j), k), and l), respectively.

distant ground user. Here, the UAV is deployed and located in the near-field region of BS while the ground users are in the far-field region of the UAV. In this case, we aim to maximize the transmission rate and minimize the power consumption by jointly optimizing the trajectory of the UAV and the power allocation of BS and UAV.

RESULT ANALYSIS

Analysis of Different Types of Action Spaces: We explore the performance of GAI-enhanced DRL in continuous action space (i.e., continuous UAV trajectory design and power allocation), discrete action space (i.e., way-point discrete UAV trajectory design and power allocation), and hybrid action spaces (i.e., discrete UAV trajectory design and continuous power allocation).

Continuous Action Space: Figure 4a, b, and c show the convergence curves of reward, transmission rate and power consumption in continuous action space, respectively. As an improved version of DDPG, TD3 significantly enhances policy stability and performance through the use of twin Q-networks and delayed policy updates. Therefore, in our case study, we select TD3 as the foundational DRL algorithm. We can observe that the GAN-enhanced TD3 algorithm shows significant improvements over the original TD3 algorithm, and this improvement is due to the novel adversarial training of the critic and target critic networks, which enhances the process of state-action value estimation. However, it is unable to effectively learn the constraints on power within the continuous solution space, resulting in penalties that cause fluctuations in the reward. Moreover, we observe that VAE improves the performance of the TD3 algorithm, which is attributed to the

capability of VAE to reduce the dimensionality of the high-dimensional state space to extract critical information. However, the process of transforming the high-dimensional state space into a low-dimensional latent space inevitably results in loss of information and leads to instability. Furthermore, we find that the Transformer-enhanced TD3 algorithm barely improves the performance of the original TD3 algorithm. This may be because Transformer is better suited for handling long-range problems, which may not align well with the current optimization problem. In addition, we observe that GDM-enhanced-TD3 achieves the best reward values and transmission rates compared to other algorithms. This is because GDM can accurately capture the underlying data distribution, providing a more effective representation of the environment. Moreover, the unique structure of GDM, which involves a diffusion process, offers a more stable and efficient learning process.

Discrete Action Space: The convergence curves of reward, transmission rate and power consumption in the discrete action space are shown in Fig. 4e, f, and g, respectively. Unlike its performance in continuous action space, GAN-enhanced GAI demonstrates superior convergence speed and quality in all of metrics above. This is because the discrete action space we designed inherently satisfies the power constraints of the optimization problem, eliminating the need for GAN to learn these constraints. However, this also indirectly highlights the weakness of GAN in learning constraint conditions directly related to optimization variables. Moreover, the convergence curves of VAE-enhanced GAI in discrete space are relatively unstable. This could be due to the fact that in discrete action spaces, the decision boundaries are

Although GAI models can improve the performance of DRL algorithms significantly, these models inevitably increase the computational complexity of the original DRL algorithms, which requires longer training time. Therefore, it is crucial to enhance the computational efficiency of GAI-enhanced DRL algorithms.

usually more complex and require precise state information to make correct action choices. However, the latent space obtained after dimensionality reduction by the VAE model may not provide sufficiently accurate information to correctly distinguish different discrete actions, leading to decision biases. Furthermore, we observe from the figure that the GDM-based algorithm still achieves the excellent performance, which indicates that it can handle discrete actions well.

Hybrid Action Space: Figure 4i, j, and k show the convergence curves of reward, transmission rate and power consumption in hybrid action space, respectively. We can observe that the GDM-enhanced TD3 is optimal in terms of the convergence speed and stability on the hybrid action space, and it converges to the best values with regard to reward, transmission rate, and power consumption. In the hybrid action space, the discrete and continuous actions are not independent, and different discrete actions correspond to different optimal continuous actions [10]. Therefore, to make the best hybrid action decisions, it is essential to explore the latent relationships between discrete and continuous actions. Evidently, the strong capability of GDM in learning the underlying data relationships enhances the accuracy and stability of TD3 decisions, enabling GDM-enhanced TD3 to achieve outstanding performance in the hybrid action space.

Analysis of Multiple GAI Models Combination: We further combine multiple GAI models in TD3 to explore their performance on the above-mentioned three action spaces. Specifically, we propose the GDM-GAN-enhanced TD3 algorithm, which introduces GAN into the GDM-enhanced TD3 algorithm to improve the critic network. Based on this, we design the GDM-GAN-Transformer-enhanced TD3 algorithm, which improves the denoising network of GDM by using Transformer. Finally, the GDM-GAN-Transformer-VAE-enhanced TD3 algorithm is proposed, which introduces VAE into the GDM-GAN-Transformer-enhanced TD3 to reduce the dimensionality of the state space. The results are shown in Fig. 4, and the corresponding results analysis is listed below.

Results of GDM-GAN-Enhanced TD3 Algorithm: As can be seen, the GDM-GAN-enhanced TD3 algorithm slightly outperforms the GDM-enhanced TD3 algorithm. This is because the GAN further enhances the quality of the estimation of state-action values by the critic network, which is beneficial for the agent to explore more effective actions.

Results of GDM-GAN-Transformer-Enhanced TD3 Algorithm: Introducing the Transformer does not lead to significant changes in the performance of the GDM-GAN-enhanced TD3 algorithm. This is because Transformer is more suitable for addressing long-range optimization problems and it does not provide substantial performance gains in our considered case study.

Results of GDM-GAN-Transformer-VAE-Enabled TD3 Algorithm: Adopting VAE results in the performance decline of the GDM-GAN-Transformer-enhanced TD3 algorithm. This decline results from the information loss caused by the dimensionality reduction of state data using VAE, which negatively impacts the GAN-based critic network that relies on detailed state information, leading to an overall performance decrease.

Training Time: The training time is also newly recorded in Fig. 4, and it can be seen from the figure that the introduced additional improved factors, that is, the GAI models, inevitably increases training time. Moreover, we can observe that the training time of GDM-enhanced TD3 is significantly higher than that of the original TD3 algorithm. This is because the denoising time step in GDM leads to an increase in the training time of the neural network, while this also leads to a notable performance improvement. Similarly, although the training time of GDM-GAN-enhanced TD3 is slightly higher than that of the GDM-enhanced TD3, the former achieves better results, which is particularly beneficial for scenarios that require high decision accuracy.

Ultimately, all the GAI models are able to improve the performance of the original TD3 algorithm for dealing with the optimization problem in the case study. Moreover, more integration combination of GAI models in DRL can be explored in the future work.

FUTURE DIRECTIONS

In this section, we present some future directions for GAI-enhanced DRL.

EFFICIENCY ENHANCEMENT OF GAI-ENHANCED DRL

Although GAI models can improve the performance of DRL algorithms significantly, these models inevitably increase the computational complexity of the original DRL algorithms, which requires longer training time. Therefore, it is crucial to enhance the computational efficiency of GAI-enhanced DRL algorithms. One approach is to decompose the complex task into multiple simpler sub-tasks and distribute these sub-tasks across multiple machines or processors during the training process. This approach enhances efficiency by focusing on smaller and more manageable issues.

GAI-ASSISTED REWARD DESIGN

As a key component of DRL, the quality of reward function design significantly impacts the performance of DRL. However, reward functions are usually manually set, and finding an appropriate reward function often requires multiple trials. In this context, the LLM integrated with retrieval-augmented generation (RAG) technique can offer a novel approach to improve the reward function design process. Specifically, the RAG model retrieves relevant information from existing reward function databases and research papers about DRL. Subsequently, the retrieved information is combined with the LLM to adjust and fine-tune the reward function for specific tasks or environments, thereby enhancing the accuracy of reward design.

GAI-IMPROVED TRANSFER LEARNING

If the application scenarios of DRL change, we usually have to retrain the DRL models from scratch, which results in a waste of computational resources. Notably, GAI shows great potential to solve this issue. Specifically, GAI can transfer the learned knowledge across different domains by identifying and exploiting common underlying structures that unify the distribution of newly generated data to a given distribution. Therefore, leveraging GAI to identify meta-knowledge that can be transferred across various domains helps bridge the gaps among different application scenarios of DRL [15].

Different GAI models possess distinct characteristics and focuses. Therefore, it is essential to explore the integration of more GAI models into DRL algorithms to address various challenges faced by different DRL algorithms and improve the decision-making process. Based on this, refining the criteria for GAI model selection can ensure that the most suitable GAI model can be quickly chosen to address specific application scenarios and challenges within DRL. Moreover, it is necessary to develop appropriate evaluation metrics to assess the impact of GAI models on the performance of DRL algorithms.

CONCLUSION

In this article, we have introduced how GAI can be applied to DRL to improve the performance. Specifically, we have first introduced the basic principles of several GAI models and DRL algorithms. Subsequently, we have discussed how GAI enhances DRL algorithms from both the data and policy perspectives. Then, we have proposed a framework to explore the improvement of DRL algorithm by four typical GAI models, namely GAN, GDM, VAE, and Transformer. Moreover, we have constructed a case study on near-field communication to validate the effectiveness of GAI-enhanced DRL. Experimental results have demonstrated that GDM shows the most significant enhancement for DRL. Finally, we have introduced four future directions to advance the wider application of GAI in DRL.

ACKNOWLEDGMENT

This research is supported by the National Natural Science Foundation of China (62272194, 62471200), the Science and Technology Development Plan Project of Jilin Province (20220101101JC), the National Research Foundation, Singapore, and Infocom Media Development Authority under its Future Communications Research & Development Programme, Defence Science Organisation (DSO) National Laboratories under the AI Singapore Programme (FCP-NTU-RG-2022-010 and FCP-ASTAR-TG-2022-003), Singapore Ministry of Education (MOE) Tier 1 (RG87/22), and the NTU Centre for Computational Technologies in Finance (NTU-CCTF).

REFERENCES

- [1] W. Shi *et al.*, "Coordinated Operation of Active Distribution Network, Networked Microgrids, and Electric Vehicle: A Multi-Agent Pp. Optimization Method," *CSEE J. Power and Energy Systems*, 2023.
- [2] A. Gupta *et al.*, "Policy-Gradient and Actor-Critic Based State Representation Learning for Safe Driving of Autonomous Vehicles," *Sensors*, vol. 20, no. 21, 2020, p. 5991.
- [3] H. Du *et al.*, "Diffusion-Based Reinforcement Learning for Edge-Enabled AI-Generated Content Services," *IEEE Trans. Mob. Comput.*, 2024.
- [4] W. Fu *et al.*, "Decision Making for Autonomous Driving via Multimodal Transformer and Deep Reinforcement Learning*," *IEEE RCAR*, 2022, pp. 481–86.
- [5] Y. Li *et al.*, "GAN-Powered Heterogeneous Multiagent Reinforcement Learning for UAV-Assisted Task Offloading," *Ad Hoc Networks*, vol. 153, 2024, p. 103341.
- [6] R. Lin *et al.*, "Physical-Layer Security Enhancement in Energy-Harvesting-Based Cognitive Internet of Things: A Gan-Powered Deep Reinforcement Learning Approach," *IEEE Internet Things J.*, vol. 11, no. 3, 2024, pp. 4899–4913.
- [7] Y. Yang *et al.*, "Decentralized Cooperative Caching and Offloading for Virtual Reality Task Based on GAN-Powered Multi-Agent Reinforcement Learning," *IEEE Trans. Serv. Comput.*, vol. 17, no. 1, 2024, pp. 291–305.

- [8] D. Chen *et al.*, "Transformer-Based Reinforcement Learning for Scalable Multi-UAV Area Coverage," *IEEE Trans. Intelligent Transportation Systems*, 2024.
- [9] H. Han *et al.*, "Primary-User-Friendly Dynamic Spectrum Anti-Jamming Access: A GAN-Enhanced Deep Reinforcement Learning Approach," *IEEE Wirel. Commun. Lett.*, vol. 11, no. 2, 2022, pp. 258–62.
- [10] B. Li *et al.*, "Hyar: Addressing Discrete-Continuous Action Reinforcement Learning via Hybrid Action Representation," *ICLR*, 2022.
- [11] X. Wang *et al.*, "Efficient Transfer Learning in Diffusion Models via Adversarial Noise," arXiv preprint arXiv: 2308.11948, 2023.
- [12] Y. Hua *et al.*, "GAN-Powered Deep Distributional Reinforcement Learning for Resource Management in Network Slicing," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 2, 2020, pp. 334–49.
- [13] H. Du *et al.*, "Beyond Deep Reinforcement Learning: A Tutorial on Generative Diffusion Models in Network Optimization," arXiv preprint arXiv:2308.05384, 2023.
- [14] C. Wu *et al.*, "Two-Stage Hierarchical Beam Training for Near-Field Communications," *IEEE Trans. Veh. Technol.*, vol. 73, no. 2, 2024, pp. 2032–44.
- [15] Y. Liu *et al.*, "Deep Generative Model and Its Applications in Efficient Wireless Network Management: A Tutorial and Case Study," *IEEE Wireless Commun.*, vol. 31, no. 4, 2024, pp. 199–207.

BIOGRAPHIES

GENG SUN (sungeng@jlu.edu.cn) received a B.S. degree in Communication Engineering from Dalian Polytechnic University, China, and the Ph.D. degree in Computer Science from Jilin University, China, in 2011 and 2018, respectively. He was a visiting researcher in the School of Electrical and Computer Engineering at Georgia Institute of Technology, USA. He is currently a Professor in the College of Computer Science and Technology at Jilin University, and his research interests include UAV Networks, collaborative beamforming, generative AI and optimizations.

WENWEN XIE (xieww22@mails.jlu.edu.cn) received the B.S. degree in Computer Science and Technology from Hefei University of Technology, Hefei, China, in 2022. She is currently working toward the MS degree in Computer Science and Technology at Jilin University, Changchun, China. Her research interests include UAV communications, IRS beamforming and deep reinforcement learning.

DUSIT NIYATO [F] (dniyato@ntu.edu.sg) is currently a professor in the School of Computer Science and Engineering, Nanyang Technological University, Singapore. He received the B.Eng. degree from King Mongkut's Institute of Technology Ladkrabang (KMUTL), Thailand, in 1999, and the Ph.D. in electrical and computer engineering from the University of Manitoba, Canada, in 2008. His research interests are in the areas of Internet of Things (IoT), machine learning, and incentive mechanism design.

FANG MEI (meifang@jlu.edu.cn) received the M.Sc. and Ph.D. degrees in computer science from Jilin University, Changchun, China, in 2005 and 2010, respectively. She is currently an Associate Professor in the College of Computer Science and Technology at Jilin University. Her research interests include intelligent information processing, multi-access edge computing, vehicular networks, and antenna arrays.

JIAWEN KANG (kavinkang@gdut.edu.cn) received the Ph.D. degree from the Guangdong University of Technology, China, in 2018. He has been a postdoc at Nanyang Technological University, Singapore, from 2018 to 2021. He is currently a full professor at Guangdong University of Technology, China. His research interests focus on blockchain, security, and privacy protection.

HONGYANG DU (duhy@hku.hk) is an assistant professor at the Department of Electrical and Electronic Engineering, The University of Hong Kong. He received his Ph.D. degree from the College of Computing and Data Science, Energy Research Institute @ NTU, Nanyang Technological University, Singapore, in 2024. He serves as the Editor-in-Chief assistant of IEEE Commun. Surveys & Tutorials (2022-2024), and the Guest Editor for IEEE Vehicular Technology Magazine. His research interests include edge intelligence, generative AI, and network management.

SHIWEN MAO [F] (smao@ieee.org) received his Ph.D. in electrical and computer engineering from Polytechnic University, Brooklyn, NY. He is a Professor and Earle C. Williams Eminent Scholar, and Director of the Wireless Engineering Research and Education Center at Auburn University. His research interests include wireless networks and multimedia communications.

Different GAI models possess distinct characteristics and focuses. Therefore, it is essential to explore the integration of more GAI models into DRL algorithms to address various challenges faced by different DRL algorithms and improve the decision-making process.