

Reinforcement Learning With LLMs Interaction for Distributed Diffusion Model Services

Hongyang Du , Ruichen Zhang , Dusit Niyato , *Fellow, IEEE*, Jiawen Kang , Zehui Xiong , Shuguang Cui , *Fellow, IEEE*, Xuemin Shen , *Fellow, IEEE*, and Dong In Kim , *Life Fellow, IEEE*

Abstract—Distributed Artificial Intelligence-Generated Content (AIGC) has attracted significant attention, but two key challenges remain: maximizing subjective Quality of Experience (QoE) and improving energy efficiency, which are particularly pronounced in widely adopted Generative Diffusion Model (GDM)-based image generation services. In this paper, we propose a novel user-centric Interactive AI (IAI) approach for service management, with a distributed GDM-based AIGC framework that emphasizes efficient and cooperative deployment. The proposed method restructures the GDM inference process by allowing users with semantically similar prompts to share parts of the denoising chain. Furthermore, to maximize the users' subjective QoE, we propose an IAI approach, i.e., Reinforcement Learning With Large Language Models Interaction (RLLI), which utilizes Large Language Model (LLM)-empowered generative agents to replicate users interactions, providing real-time and subjective QoE feedback aligned with diverse user personalities. Lastly, we present the GDM-based Deep Deterministic Policy Gradient (G-DDPG) algorithm, adapted

to the proposed RLLI framework, to allocate communication and computing resources effectively while accounting for subjective user traits and dynamic wireless conditions. Simulation results demonstrate that G-DDPG improves total QoE by 15% compared with the standard DDPG algorithm.

Index Terms—AI-generated content, reinforcement learning, generative diffusion model, generative agents, large language models.

I. INTRODUCTION

THE growing demand for Artificial Intelligence-Generated Content (AIGC) services in multimedia and business applications [1] is driven by the advancement of Generative AI (GenAI) models, which provide scalable and consistent outputs in both text and imagery [2]. ChatGPT, for example, reached over 100 million active users within two months [3], demonstrating its influence on text-based interactions. In the realm of visual content generation, Stable Diffusion's capacity to create images from text prompts shows significant progress in multi-modal technologies [4], [5]. The widespread adoption of AIGC services in human societies indicates a notable transition to Interactive AI (IAI) as the next evolutionary phase of GenAI [6]. This shift is reshaping human-content interaction and reflects ongoing developments in human-AI collaboration.

However, Quality of Experience (QoE) maximization in AIGC services emerges as a critical challenge, due to the subjective nature of human perception, which extends beyond objective image quality metrics [7], [8]. Four images in Fig. 1 exemplify this complexity, showing four distinct images that are generated under a same prompt "A man sits in the street". Although each image is of high quality, they align differently with user personality profiles, resulting in varying QoE evaluations. Therefore, network designers aim to develop service and management models that adapt to individual personalities to improve QoE. However, the lack of a precise mathematical representation of QoE makes this optimization challenging. While some studies have adopted psychological laws to approximate users' subjective QoE [9], these methods often fail to capture the complexity of real-world scenarios. An alternative is to train management models using Reinforcement Learning with Human Feedback (RLHF), which relies on expert-provided QoE labels. This method is expensive, ethically sensitive, and difficult to implement in real-time, raising the first research question:

Received 18 November 2023; revised 10 April 2025; accepted 17 June 2025. Date of publication 30 June 2025; date of current version 5 September 2025. This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant U22A2054, in part by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Future Communications Research & Development Programme under Grant FCP-NTU-RG-2022-010 and Grant FCP-ASTAR-TG-2022-003, in part by Singapore Ministry of Education (MOE) Tier 1 under Grant RG87/22 and Grant RG24/24, in part by the NTU Centre for Computational Technologies in Finance (NTU-CCTF), in part by the RIE2025 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) under Grant I2301E0026, administered by A*STAR, Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL), in part by the MSIT (Ministry of Science and ICT), Korea, under the ICT Creative Consilience program under Grant IITP-2020-0-01821, and in part by the ITRC support program under Grant IITP-2023-RS-2023-00258639 supervised by the IITP (Institute for ICT Planning & Evaluation). Recommended for acceptance by Y.-Z. Song. (*Corresponding author: Hongyang Du.*)

Hongyang Du is with the Department of Electrical and Electronic Engineering, University of Hong Kong, Hong Kong, SAR, China (e-mail: duhy@eee.hku.hk).

Ruichen Zhang and Dusit Niyato are with the College of Computing and Data Science, Nanyang Technological University, Singapore 639798 (e-mail: ruichen.zhang@ntu.edu.sg; dnyato@ntu.edu.sg).

Jiawen Kang is with the School of Automation, Guangdong University of Technology, Guangzhou 510006, China (e-mail: kavinkang@gdut.edu.cn).

Zehui Xiong is with the School of Electronics, Electrical Engineering and Computer Science (EEECSS), Queen's University Belfast, BT7 1NN Belfast, U.K. (e-mail: z.xiong@qub.ac.uk).

Shuguang Cui is with the School of Science and Engineering (SSE) and the Future Network of Intelligence Institute (FNii), Chinese University of Hong Kong (Shenzhen), Shenzhen 518172, China (e-mail: shuguangcui@cuhk.edu.cn).

Xuemin Shen is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: sshen@uwaterloo.ca).

Dong In Kim is with the Department of Electrical and Computer Engineering, Sungkyunkwan University, Suwon 16419, South Korea (e-mail: dongin@skku.edu).

Digital Object Identifier 10.1109/TPAMI.2025.3584698

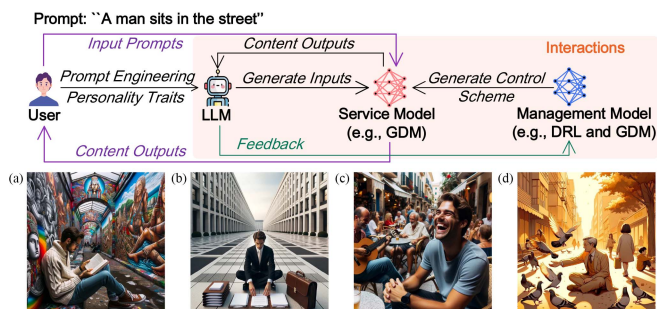


Fig. 1. The basic framework of IAI and four images generated with the prompt “A man sits in the street”. Part A is a man engrossed in a book against vibrant street art appeals to users with high *openness*. Part B is a formally dressed man on a clean street, resonating with users high in *conscientiousness*. Part C is a sociable street cafe scene, suitable for users with high *extraversion*. Part D is a man feeding pigeons in a peaceful setting with playing children, fitting for users with high *agreeableness*.

C1. How to obtain human-aware subjective QoE feedback efficiently and design the communication and computing resource allocation algorithm?

Addressing *CI* requires the exploration of cutting-edge approaches, among which IAI stands out as a promising solution [6]. IAI focuses on designing AI models that learn and adapt through user interaction, progressively advancing AI models’ performance and operational efficacy. This paradigm shift, from static to dynamic learning systems, equips IAI-enhanced networks with the ability to offer tailored responses and proactively adapt, marking a significant stride in personalized AI-based service managements. For AIGC services, as shown in Fig. 1, we propose Reinforcement Learning With LLM Interaction (RLLI) algorithm as one step towards IAI, incorporates Deep Reinforcement Learning (DRL) for its suitability in dynamic environments [10] and Large Language Model (LLM) for advanced knowledge understanding and generative capabilities:

- DRL regards user QoE as quantifiable rewards and circumvents complexity of mathematically modeling subjective QoE. When integrated with GenAI techniques such as Generative Diffusion Model (GDM) [11] that effectively handle complex data modeling and produce high-quality samples, DRL can be further enhanced to allocate network resources efficiently while factoring in user satisfaction.
- LLM-empowered Generative Agents (GAs) can represent real AIGC users with various personalities to generate QoE feedback, minimizing human resource input and associated ethical risks. By embedding human personality traits into GAs through prompts [12], these GAs can simulate human interaction in the training of DRL algorithms.

Furthermore, AIGC models impose substantial energy demands during the inference phase [5]. Particularly in GDM-based image generation, multiple denoising steps are essential for producing high-quality outputs. This iterative process progressively removes noise, resulting in considerable energy and time consumption [5]. Although few-step models have been proposed to reduce computation and accelerate inference [13], [14], they often degrade image quality and detail compared to

multi-step approaches [15], [16]. In practice, widely used AIGC systems such as Stable Diffusion [4] and Sora [17] continue to depend on multi-step denoising, underscoring its necessity for achieving high-quality results. Additionally, transmitting these high-quality images, particularly in latency-sensitive applications like augmented reality and telemedicine, requires significant network bandwidth [18], placing increased demands on future Sixth-Generation (6 G) networks that aim to support human-centric and high QoE services [19]. These challenges lead to the second research question:

C2. How to harness network capabilities to facilitate energy-efficient and low-latency GDM inference?

Fortunately, advancements in model compression facilitate the deployment of large GenAI models directly on network edge devices [20], [21]. This forms a foundation for network-assisted AIGC, particularly collaborative distributed diffusion model [22], to address *CI*. As shown in Part A of Fig. 2, by distributing the denoising steps across multiple network edge devices, we can achieve a more flexible and efficient network resource utilization. More importantly, generation tasks with semantically similar prompts can share several denoising steps, thereby saving overall energy and time consumption. As shown in Part B of Fig. 2, the denoising *Path 2*, which starts with “Prompt 2: An apple on table” and switches to “Prompt 1: A lemon on table” can still converge nearly with *Path 1* that uses “Prompt 1” throughout. We can observe from Part C of Fig. 2 that when employing an appropriate distributed denoising process, such as switching prompts in an early stage or maintaining semantic similarity between prompts, high-quality images can be generated that is corresponding to user prompts while reducing overall resource consumption. For example, we can observe a clear reduction in energy and time costs when generating images via *Path 2* and *Path 3*, compared to independent operations through *Path 1* and *Path 3*.

Note that the distributed GDM-based AIGC framework enables semantic fusion of prompts from multiple users, making the quality of generated images dependent on the designed communication and computing resource allocation scheme. This dependency is addressed by the proposed edge-based IAI solution, which introduces a new optimization dimension to enhance AIGC service quality and improve user QoE. The contributions of this paper are summarized as follows.

- We introduce a collaborative distributed GDM-based AIGC framework for image generation services. This approach enhances user privacy by generating the final content on edge devices, reducing the total network energy and time overheads (for *CI*).
- We present a QoE feedback scheme by using LLM-empowered GAs to simulate the human’s different personalities. With the aid of prompts and assigning one agent per user, GAs can mimic users of diverse subjective preferences, delivering evaluations of quality of generated images (for *C2*).
- We propose an IAI algorithm, i.e., RLLI, with LLM-generated QoE at its core. The goal is to determine the optimal denoising steps, i.e., computing energy consumption, at server and user devices, and downlink transmission

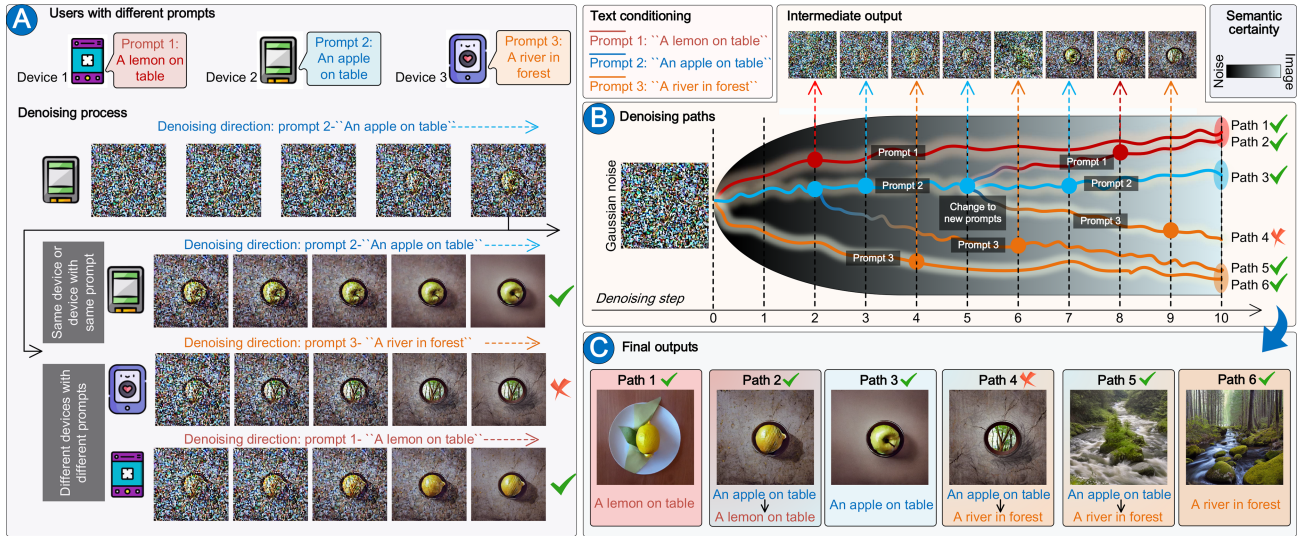


Fig. 2. The working principle of the GDM and motivations behind distributed denoising inference process. *Part A* depicts the cooperative inference process across devices where, starting with Gaussian noise on *Device 2*, it denoises using *Prompt 2* before *Devices 1* and *3* continue in succession towards their respective prompts. *Part B* shows the fundamentals of distributed GDM inference, illustrating how denoising path directions alter with changing prompts, and emphasizing the high semantic similarity of *Prompt 1* to *Prompt 2*, contrasting it with the distance between *Prompt 2* and *Prompt 3*. Consequently, *Path 2* aligns with *Prompt 1*'s requirements, unlike *Path 4* with *Prompt 3*. *Part C* showcases the final generation outcomes for all five paths.

power, i.e., communications energy consumption, of the server to user devices, based on the user personalities and wireless conditions (for *C2*).

The rest of this paper is organized as follows: Section II reviews related literature. In Section III, we describe the system model and formulate the problem for our distributed GDM-based AIGC framework. Section IV details the RLLI algorithm. Section V conducts an evaluation of the proposed framework and RLLI algorithm. The paper concludes with Section VI. A list of mathematical symbols and functions most frequently used in this paper is available in Table I, while Table II summarizes the abbreviations.

II. RELATED WORK

In this section, we discuss several related works, including GDM, aesthetics analysis, and LLM.

A. Generative Diffusion Models

GDMs are essential to various AIGC services. Their applications include image and video generation in computer vision using DDPM [23] and video diffusion models [24], text generation and editing through Diffusion-LM [25] and DiffuSER [26], audio synthesis with ProDiff [27] and DiffWave [28], as well as specialized tasks such as graph generation [29] and molecular structure creation [30]. As privacy-sensitive applications grow, there is increasing interest in deploying GDMs on edge devices. While model compression has made this feasible, the inference process—particularly the iterative denoising steps—remains energy-intensive, increasing the total energy cost across the network. However, the sequential structure of the denoising chain enables distributed computation. Cooperative distributed GDM [22] exploits this structure by allowing devices to share

TABLE I
MATHEMATICAL NOTATIONS

c_k	The prompt from the k -th user, where $k = 1, \dots, K$
\mathcal{D}	The decoder that decodes a latent representation in the latent space into an image in the RGB space
δ_0	The energy cost for each denoising step of the server
δ_k	The energy cost for each denoising step of the k -th device
E_T	The energy budget in the server
E_{T_k}	The energy budget in the k -th device
\mathcal{E}	The encoder that encodes an image in the RGB space into a latent representation in the latent space
g_k	The wireless channel gains from the server to the k -th device
K	The total number of users
P_k	The transmission power from the server to the k -th device
Q_k	The QoE of the k -th user
Q_{th}	The QoE threshold
t_0	The number of denoising steps in the server
t_k	The number of denoising steps in the k -th device
$\tau(\cdot)$	The domain-specific encoder that projects c to a latent representation
u_k	The personality of the k -th user
ϵ_θ	The denoising network of the GDM

partial denoising steps, which reduces redundant computation and lowers network-wide energy consumption.

B. Aesthetics Analysis and Big Five Personality Traits

Traditional approaches to Image Aesthetic Assessment (IAA) have focused on universal aesthetic scores, overlooking the subjective nature of aesthetic preferences driven by individual personality differences. Recent developments, such as those

TABLE II
LIST OF ABBREVIATIONS

Abbreviation	Full Term
AIGC	Artificial Intelligence-Generated Content
QoE	Quality of Experience
GDM	Generative Diffusion Model
IAI	Interactive AI
VIT	Visual Instruction Tuning
LLM	Large Language Model
RLLI	Reinforcement Learning With LLM Interaction
DRL	Deep Reinforcement Learning
RLHF	Reinforcement Learning with Human Feedback
DDPG	Deep Deterministic Policy Gradient
G-DDPG	GDM-based Deep Deterministic Policy Gradient
GA	Generative Agent
BEP	Bit Error Probability
SNR	Signal-to-Noise Ratio

integrating the Big Five personality traits into a multi-task deep learning model [31] for IAA, represent a paradigm shift. The Big Five personality traits [32], comprising openness, conscientiousness, extraversion, agreeableness, and neuroticism, provide a widely accepted structure for characterizing personality. Specifically, openness captures creativity and a willingness of users to explore new experiences; conscientiousness involves diligence and reliability; extraversion denotes sociability and assertiveness; agreeableness reflects cooperativeness and kindness; and neuroticism is associated with emotional instability and anxiety. Embedding these traits into IAA enables a transition from general scoring to individualized assessments, offering the potential to align AIGC models with users' psychological profiles and make AIGC services more personalized [33].

C. LLM-Empowered Generative Agents as Evaluators

LLMs-empowered GAs can interpret and respond to complex natural language instructions, making them a versatile interface for a range of tasks, including evaluative ones [34], [35]. Specifically, the authors in [34] explore the capabilities of ChatGPT in evaluating textual content across various human-aware criteria such as quality, tone, and coherence. These evaluations provide a foundation for extending LLM functionality into other domains. Subsequently, the authors in [36] evaluate the potential of LLMs like ChatGPT in computational social science, examining their ability to perform both classification and generative tasks in a zero-shot manner. It is shown that LLMs demonstrate fair agreement with humans and can augment the annotation process. More recently, the authors in [37] show that modern role-playing LLMs can successfully mimic specific personality traits, achieving an 82.8% alignment with human perceptions. These findings support the use of LLM-empowered GAs for simulating users with diverse subjective preferences in image evaluation tasks, positioning GA feedback as a viable alternative to human feedback for training service management algorithms.

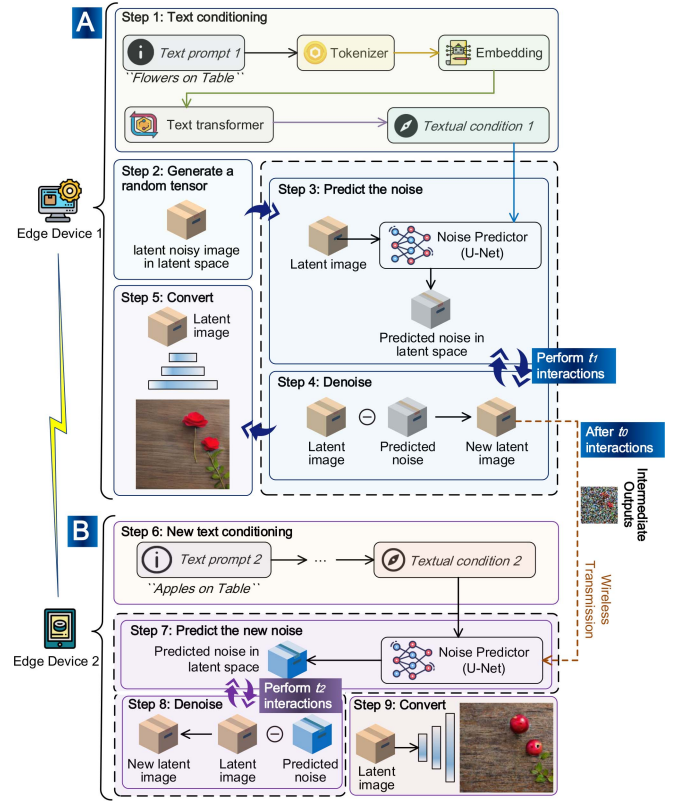


Fig. 3. The working steps of GDM-based text-to-image generation service, and the principle to achieve multi-device distributed denoising process.

III. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we analyze the principles and conditions under which GDMs can be executed separately on different devices. Furthermore, we discuss the deployment method and formulate the user-centric QoE optimization problem in our proposed distributed GDM-based AIGC framework.

A. Distributed Generative Diffusion Model-Based AIGC

1) *Latent Diffusion Models*: GDMs, as illustrated in Part A of Fig. 3, are probabilistic models that approximate a data distribution by iteratively denoising samples drawn from a normal distribution, effectively reversing a Markov chain over T steps. The Stable Diffusion [4] stands out by adopting latent GDM, operating in the latent space rather than the high-dimensional image pixel space. Instead of manifesting a perturbed image, latent GDM introduces “latent noise”, i.e., a stochastic tensor in the latent space, during training to perturb the latent representation. For example, an image with 512×512 resolution, the corresponding latent size is a much smaller, i.e., $4 \times 64 \times 64$, enabling faster processing and less memory usage.

Diving deeper into the mechanics, we consider $\mathbf{x}_0 \in \mathbb{R}^{H \times W \times 3}$ is an image in Red, Green & Blue (RGB) space. The encoder \mathcal{E} encodes \mathbf{x}_0 into a latent representation as

$$\mathbf{z}_0 = \mathcal{E}(\mathbf{x}_0), \quad (1)$$

where $\mathbf{z}_0 \in R^{h \times w \times r}$. The encoder \mathcal{E} downsamples the image \mathbf{x}_0 by a factor f , i.e., $f = H/h = W/w$. The decoder \mathcal{D} reconstructs the image from the latent representation as

$$\tilde{\mathbf{x}}_0 = \mathcal{D}(\mathbf{z}_0) = \mathcal{D}(\mathcal{E}(\mathbf{x}_0)). \quad (2)$$

GDMs can be specified in terms of a Signal-to-Noise Ratio (SNR) as

$$\gamma(t) = \frac{\alpha_t^2}{\sigma_t^2}, \quad (3)$$

where α_t and σ_t are strictly positive scalar-value functions of t , and $\gamma(t)$ is strictly monotonically decreasing in time, i.e., $\gamma(t) < \gamma(s)$ for any $t > s$. We consider that both α_t and σ_t are smooth, such that their derivatives with respect to t are finite [38]. This diffusion process specification includes the variance-preserving diffusion process as a special case [23], [38], where $\alpha_t = \sqrt{1 - \sigma_t^2}$. Another special case is the variance-exploding diffusion process [39], [40], where $\alpha_t^2 = 1$. We use the variance-preserving version in this paper. By starting from a data sample \mathbf{z}_0 , a forward diffusion process q can be defined as

$$q(\mathbf{z}_t | \mathbf{z}_0) = \mathcal{N}(\mathbf{z}_t | \alpha_t \mathbf{z}_0, \sigma_t^2 \mathbf{I}). \quad (4)$$

The Markov structure for $t > s$ can be expressed as

$$q(\mathbf{z}_t | \mathbf{z}_s) = \mathcal{N}(\mathbf{z}_t | \alpha_{t|s} \mathbf{z}_s, \sigma_{t|s}^2 \mathbf{I}), \quad (5)$$

where $\alpha_{t|s} = \alpha_t / \alpha_s$ and $\sigma_{t|s}^2 = \sigma_t^2 - \alpha_{t|s}^2 \sigma_s^2$.

Denoising GDMs are generative models p which revert the forward diffusion process with a similar Markov structure running backward as

$$p(\mathbf{z}_0) = p(\mathbf{z}_T) \prod_{t=1}^T p(\mathbf{z}_{t-1} | \mathbf{z}_t). \quad (6)$$

The prior $p(\mathbf{z}_T)$ is typically chosen as a standard normal distribution. We parameterize $p(\mathbf{z}_{t-1} | \mathbf{z}_t)$ to specify it in terms of the true posterior $q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0)$. We use the network with parameter θ , i.e., $\varepsilon_\theta(\mathbf{z}_t, t)$, to estimate the unknown original data sample \mathbf{z}_0 . Then, we obtain that [41]

$$\begin{aligned} p(\mathbf{z}_{t-1} | \mathbf{z}_t) &= q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) \\ &= q(\mathbf{z}_{t-1} | \mathbf{z}_t, \varepsilon_\theta(\mathbf{z}_t, t)) \\ &= \mathcal{N}\left(\mathbf{z}_{t-1} | \mu_\theta(\mathbf{z}_t, t), \sigma_{t|t-1}^2 \frac{\sigma_{t-1}^2}{\sigma_t^2} \mathbf{I}\right), \end{aligned} \quad (7)$$

where

$$\mu_\theta(\mathbf{z}_t, t) = \frac{\alpha_{t|t-1} \sigma_{t-1}^2}{\sigma_t^2} \mathbf{z}_t + \frac{\alpha_{t-1} \sigma_{t|t-1}^2}{\sigma_t^2} \varepsilon_\theta(\mathbf{z}_t, t). \quad (8)$$

We use the reparameterization [23]

$$\varepsilon_\theta(\mathbf{z}_t, t) = \frac{\mathbf{z}_t - \alpha_t \mathbf{z}_\theta(\varepsilon_t, t)}{\sigma_t}, \quad (9)$$

to express the reconstruction term as a denoising objective.

For convenience, we denote the denoising process that involves (7), (8), and (9) as $\text{Denoise}(\cdot)$

$$\mathbf{z}_{t-1} = \text{Denoise}(\mathbf{z}_t, t). \quad (10)$$

Algorithm 1: Distributed Latent GDM Inference Process Among K Devices.

Input: the textual prompt $\{\mathbf{c}_1, \dots, \mathbf{c}_K\}$ from devices $1, \dots, K$, respectively

Output: Generated image $\tilde{\mathbf{x}}_0$ for any device k ($k = 2, \dots, K$)

```

1: ## In Device 1
2: Sample initial Gaussian noise  $\mathbf{z}_T$  in the latent space
3: for  $t = T$  to  $t_{\text{off}} + 1$  do
4:    $\mathbf{z}_{t-1} \leftarrow \text{Denoise}(\mathbf{z}_t, t, \tau(\mathbf{c}_1))$ , i.e., (13)
5: end for
6: ## Wireless Transmission
7: Transmit the  $\mathbf{z}_{t_{\text{off}}}$  from device 1 to devices  $2, \dots, K$ 
8: ## In Device  $k$  ( $k = 2, \dots, K$ )
9: for  $t = t_{\text{off}}$  to 1 do
10:    $\mathbf{z}_{t-1} \leftarrow \text{Denoise}(\mathbf{z}_t, t, \tau(\mathbf{c}_k))$ , i.e., (13)
11: end for
12: ## Image Reconstruction
13:  $\tilde{\mathbf{x}}_0 \leftarrow \mathcal{D}(\mathbf{z}_0)$ 
14: return  $\tilde{\mathbf{x}}_0$ 

```

The corresponding loss function to train the latent GDM can be simplified to

$$\mathcal{L}_{GDM} = E_{\mathbf{z}, \varepsilon \sim \mathcal{N}(0,1), t} \left[\|\varepsilon - \varepsilon_\theta(\mathbf{z}_t, t)\|_2^2 \right], \quad (11)$$

with t uniformly sampled from $\{1, \dots, T\}$. The neural backbone is realized as a time-conditional UNet [42]. Since the forward process is fixed, \mathbf{z}_t can be efficiently obtained from \mathcal{E} during training, and samples from $p(\mathbf{z})$ can be decoded to image space with a single pass through \mathcal{D} .

2) *Conditioning Mechanisms:* Similar to other types of generative models [43], [44], GDMs are capable of modeling conditional distributions. This can be implemented with a conditional denoising auto-encoder $\varepsilon_\theta(\mathbf{z}_t, t, \mathbf{c})$ to control the synthesis process through inputs \mathbf{c} such as text [45], semantic maps [46] or other image-to-image translation tasks [47]. To pre-process \mathbf{c} from various modalities, such as language prompts, a domain-specific encoder is introduced that projects \mathbf{c} to a latent representation $\tau(\mathbf{c})$. Based on image-conditioning pairs, we then train the conditional latent GDM via

$$\mathcal{L}_{LDM} = E_{\mathbf{z}, \varepsilon \sim \mathcal{N}(0,1), t, \mathbf{c}} \left[\|\varepsilon - \varepsilon_\theta(\mathbf{z}_t, t, \tau(\mathbf{c}))\|_2^2 \right]. \quad (12)$$

In the inference process, the denoising functions can be updated accordingly, by adding \mathbf{c} as the condition together with step information t as

$$\mathbf{z}_{t-1} = \text{Denoise}(\mathbf{z}_t, t, \tau(\mathbf{c})). \quad (13)$$

3) *Distributed Diffusion Model Inference:* The distributed GDM inference process relies on sequential denoising guided by textual prompts. We analyze mathematically the conditions for changing textual prompts while still ensuring that the denoising results remain reliable. Specifically, we consider the latent space, in which the diffusion process takes place, to be the euclidean space denoted as \mathcal{L} . Every point in \mathcal{L} corresponds to a latent vector. The analysis includes the following steps:

- 1) *Defining Denoising Path and Semantic Distance*: For each textual prompt c , the denoising path, evolving from the initial Gaussian noise to the targeted semantic representation, is represented by $P_c : [T, 0] \rightarrow \mathcal{L}$. Here, $P_c(T)$ represents the Gaussian noise point, i.e., the start point in Part B of Fig. 2, while $P_c(0)$ corresponds to the latent representation of the final generated image under the prompt c , i.e., the end point in Part B of Fig. 2.

The semantic distance between two prompts, c_i and c_j , in the latent space is defined by the euclidean distance between their respective denoising path endpoints as

$$d(c_i, 0, c_j, 0) = \|P_{c_i}(0) - P_{c_j}(0)\|. \quad (14)$$

The semantic distance provides a direct measure: if the semantic information of c_i and c_j are closely related, their corresponding semantic representations, i.e., latent vectors, and the denoising paths under their guidance, should also be proximate in \mathcal{L} . Consequently, a smaller $d(c_i, 0, c_j, 0)$ indicates that substituting c_i for c_j is less likely to degrade the quality of generated images.

- 2) *Semantic Analysis*: Another key observation is that extending the denoising path, i.e., $P_c(t)$ with t from T to 0 , progressively incorporates the semantic information of the textual prompt c into the path. To measure the difference in semantic information encoded within the latent vectors at step t between prompts c_i and c_j , we employ the Kullback-Leibler (KL) divergence:

$$\mathcal{K}(p(z_t|c_i) || p(z_t|c_j)) = \sum_{z \in \mathcal{Z}_t} p(z|c_i) \log \frac{p(z|c_i)}{p(z|c_j)}. \quad (15)$$

At the beginning of the denoising process, i.e., $t = T$, the noise component dominates, leading to a latent space with high entropy due to the absence of any meaningful semantic information. Thus, both prompts would produce almost similar distributions for the latent vectors. As a result, the KL divergence, \mathcal{K} , would be small, suggesting minimal difference between the paths P_{c_i} and P_{c_j} , which means that $d(c_i, T, c_j, T)$ is small for any c_i and c_j . With the progression of denoising, the influence of the prompts begins to introduce semantic information, leading to a reduction in entropy.

- 3) *Switching Denoising Path*: Assuming a denoising path have been progressed t_{off} steps under c_i to an intermediate point $P_{c_i}(t_{\text{off}})$, where $0 < t_{\text{off}} < T$. Now, the flowing denoising path should cover the semantic distance:

$$d(c_i, t_{\text{off}}, c_i, 0) = \|P_{c_i}(t_{\text{off}}) - P_{c_i}(0)\|. \quad (16)$$

However, if we intend to switch the prompt from c_i to c_j , the distance from $P_{c_i}(t_{\text{off}})$ to the end-point of c_j is:

$$d(c_i, t_{\text{off}}, c_j, 0) = \|P_{c_i}(t_{\text{off}}) - P_{c_j}(0)\|. \quad (17)$$

This distance $d(c_i, t_{\text{off}}, c_j, 0)$ is crucial as it informs the extent of path correction needed to align the denoising process towards the semantics of c_j . From the above analysis, we conclude the following conditions for our distributed GDM-based AIGC framework:

- The new prompt c_j should have close semantic information related to c_i , reducing the difference between (16) and (17).
 - The transition to c_j should be made in a relatively early stage of the denoising process, i.e., when \mathcal{K} as given in (15) is small, to reduce the influence of c_i and allowing c_j to predominantly lead the denoising direction.
- 4) *Image Reconstruction*: The image \tilde{x}_0 reconstructed from the latent representation $P_{c_j}(0)$ reflects the semantic influence of c_j :

$$\tilde{x}_0 \approx \mathcal{D}(P_{c_j}(0)). \quad (18)$$

The decoder \mathcal{D} maps the latent representation back to the image space. A high quality in \tilde{x}_0 means the generated image is a precise representation, effectively capturing the semantics associated with c_j .

The overall algorithm for distributed latent GDM inference process among K devices is given as Algorithm 1. Note that our framework efficiently supports both synchronous and asynchronous operations. In synchronous operations, intermediate denoised results are directly transmitted to edge users without needing storage. In asynchronous operations, we can use a caching mechanism to store denoised results, as discussed in Section VI.

B. Deployment Method

The proposed distributed GDM-based AIGC framework leverages the flexible structure of the denoising process. By sharing intermediate denoised results across network devices, it reduces the system's overall energy consumption. Three network architectures are considered for deployment:

- 1) The *Edge-to-Multiple Devices* architecture executes shared denoising steps at an edge server, and then transmitting intermediate denoised results to the user devices, as shown in Fig. 4.
- 2) The *Device-to-Device* approach fosters a partnership between two devices, jointly executing denoising tasks, optimizing energy usage, and safeguarding user privacy, as shown in Fig. 3.
- 3) The strategy of *Forming a Cluster among Multiple Devices*, either with the assistance of an edge server or through self-organization, bolsters system adaptability and scalability to local requirements and environments, as shown in Part A of Fig. 2.

These architectures introduce additional design considerations under real-world conditions, such as network variability and device heterogeneity. Transmission delays, errors, and device constraints, e.g., memory, should be accounted for to ensure stable performance. Our current analysis partially reflects these through channel modeling and energy cost formulations, but further system-level adaptations could be integrated as extensions. Without loss of generality, we focus on the *Edge-to-Multiple Devices* architecture for the distributed GDM-based AIGC framework.

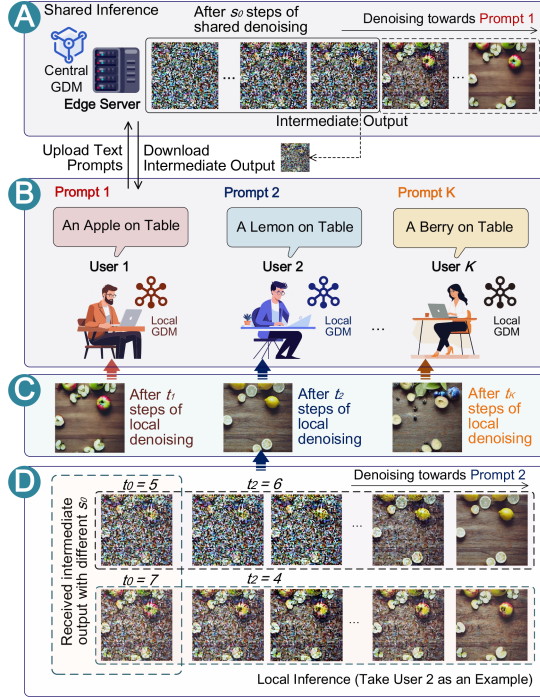


Fig. 4. Deployment method for distributed GDM-based text-to-image AIGC services: *Part A* illustrates the shared inference mechanism where users with semantically similar prompts collaborate with the server. *Part B* highlights the diverse text-to-image requirements of each of the K concurrent users. Following initial server processing, *Part D* delves into the distributed denoising process by each individual user. *Part C* presents the final generated images corresponding to user-specific prompts.

Here, we consider a server executes t_0 denoising steps, followed by K users, each executing t_k steps. Each server-side denoising step incurs an energy cost δ_0 , while every user device has its unit energy cost δ_k . Additionally, transmitting the intermediate result to the k -th user requires transmission power P_k , with a unit energy cost β . Transmitted the intermediate result to the k -th user requires transmission power P_k with unit energy cost β . The QoE depends on both the allocation of communication and computing resources and the user's personality \mathbf{u}_k , which reflects aesthetic preferences and leads to subjective QoE evaluations. Here, the resource allocation variables for communication and computing are

$$\mathbf{t} = \{t_0, t_1, \dots, t_K\}, \quad (19)$$

and

$$\mathbf{P} = \{P_1, \dots, P_K\}. \quad (20)$$

The resource allocation variables, \mathbf{t} and \mathbf{P} , are significant in determining the efficacy of the AIGC service:

- The transmission power, \mathbf{P} , impacts the bit error probability (BEP) during wireless transmission. Lower power levels risk degrading the integrity of the transmitted intermediate denoised results, potentially diminishing the quality of generated images.
- The number of denoising steps, \mathbf{t} , directly affects the quality of generated images [5]. Specifically, the shared

steps, t_0 , set the initial latent representation. We dynamically and optimally adjust t_0 to balance efficiency and personalization, enhancing QoE while managing energy consumption.

- Notably, when transmission power is insufficient, which increases the bit errors in intermediate denoised results, adding more denoising steps can improve the final image quality [22]. These additional steps function as an error-correction phase, reducing inaccuracies and improving the final result.

C. User-Centric QoE Maximization Problem

The central challenge is maximizing subjective QoE across all users while considering user personalities, objective wireless conditions, energy constraints, and QoE thresholds. This integrated modeling approach captures two essential aspects of distributed AIGC services. First, using QoE as a utility function enables personalized content generation that aligns with user expectations. Second, incorporating energy constraints is necessary for efficient energy management across devices and network nodes. Mathematically, the objective and constraints are expressed as:

$$\max_{\{\mathbf{t}, \mathbf{P}\}} \sum_{k=1}^K Q_k(t_0, t_k, P_k, \mathbf{u}_k), \quad (21)$$

$$\text{s.t.}, \quad \delta_k t_k \leq E_{T_k}, \quad \forall k \in \{1, \dots, K\} \quad (22)$$

$$Q_k(t_0, t_k, P_k, \mathbf{u}_k) \geq Q_{th}, \quad \forall k \in \{1, \dots, K\} \quad (23)$$

$$\delta_0 t_0 + \beta \sum_{k=1}^K P_k \leq E_T. \quad (24)$$

The first constraint reflects the computing energy limitations of edge devices in executing denoising steps. The second constraint ensures that the AIGC service satisfies users' minimum QoE requirements. The third constraint, which is the most critical, addresses the overall energy budget and requires a trade-off between communication and computation. Specifically, increasing transmission power reduces the BEP, thereby improving the accuracy of transmitted intermediate denoised results. However, excessive transmission power reduces the energy available for performing denoising steps, which negatively impacts the final image quality. Conversely, if the transmission power is insufficient, even enhanced denoising at the edge device cannot fully correct the errors introduced during transmission.

In the following, we use the RLLI method for the QoE maximization problem in our proposed distributed GDM-based AIGC framework.

IV. REINFORCEMENT LEARNING WITH LLMs FEEDBACK

In this section, we present the RLLI method for evaluating the QoE of the distributed GDM framework described in Section III. We first introduce an aesthetic-aware QoE model, then explain how LLM-empowered GAs provide subjective QoE feedback, which serves as the reward signal in DRL algorithms.

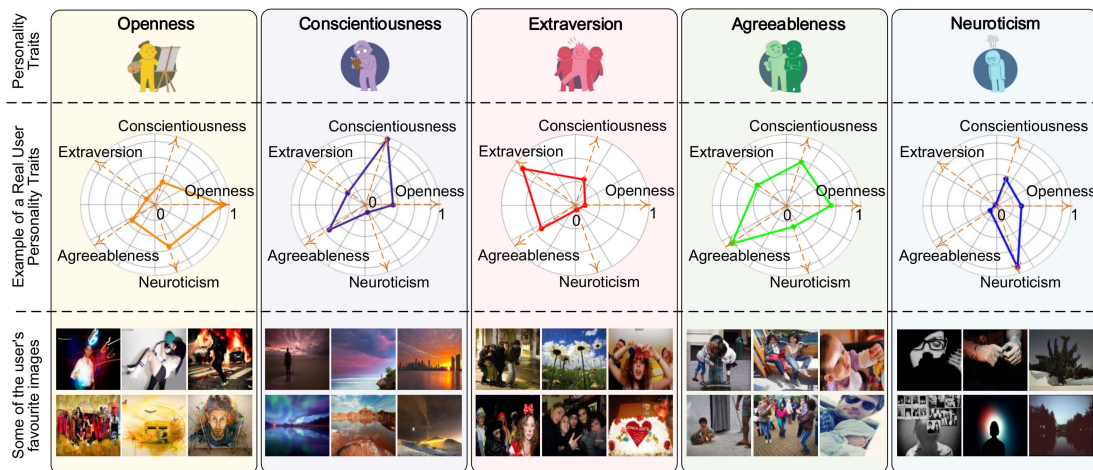


Fig. 5. Five types of personality traits in the Big-Five personality model: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism. We show real user personality scores from the PsychoFlickr database, along with examples of the images that they prefer.

A. Aesthetic-Aware QoE Modeling

1) *Big Five Personality Traits*: Incorporating GAs into the DRL training process enables the integration of human subjective factors, particularly aesthetic preferences, which are critical to QoE in image-related AIGC services. Prior studies indicate that aesthetic preferences are influenced by individual personality traits [31], [32], [33]. The Big Five personality traits [32], also known as the five-factor model, provide a widely accepted framework for characterizing personality differences.¹ The Big Five model consists of five broad dimensions of personality traits including

- *Openness*: A trait characterized by a deep affinity for imagination, novel experiences, and a broad spectrum of interests.
- *Conscientiousness*: Denotes meticulousness and structure, often manifesting in methodical, goal-oriented behaviors.
- *Extraversion*: Embodies individuals who thrive in social interactions, drawing energy from a company of others.
- *Agreeableness*: Reflects proclivities towards trustworthiness, altruism, and prosocial behaviors.
- *Neuroticism*: Typified by emotional fluctuations and heightened sensitivity to environmental stressors and adversities.

Each of these traits represents a continuous spectrum of personality characteristics. Considering that recent studies have shown LLMs can effectively simulate the Big Five personality traits with an 82.8% alignment to human perception [37], we adopt this model to guide LLM-empowered GAs in mimicking AIGC service users' personalities [12]. However, while personality traits provide a structured foundation for modeling user preferences, actual behavior may still vary in real-world contexts. We do not explicitly model such dynamics in this paper, but this represents a valuable direction for future refinement.

In Fig. 5, we display five exemplary user profiles from the PsychoFlickr dataset [33], showing their Big Five

¹One example of open source Big Five personality traits tests: <https://bigfive-test.com>.

personality scores and selected preferred images. These profiles offer empirical evidence of the correlation between individual personality traits and aesthetic preferences in image selection. This connection becomes especially important in the context of shared denoising within distributed GDMs. Unlike independent denoising, shared denoising allows semantically similar prompts to reuse intermediate diffusion steps, which improves efficiency but introduces interdependencies among outputs. As discussed in Section III-A3, the more steps that are shared, the more likely it is that the final images will influence one another. From a QoE perspective, this creates a potential conflict: users with distinct aesthetic preferences—shaped by their personalities—may experience degraded satisfaction if shared denoising alters the semantic or stylistic fidelity of their expected outputs. For instance, as depicted in Part B of Fig. 6, one user, scoring high in agreeableness, might prefer an image generated from the prompt “A dog sits in front of a bush” that has a well-groomed dog. However, if this prompt shares several shared denoising steps with another prompt such as *A tiger sits in front of a bush*, the final generated images for this user may have a dog that exhibits stylistic features of a tiger, causing the low QoE. Thus, the settings for shared denoising steps are pivotal in achieving maximum sum QoE.

2) *Prompt Design for LLM-Empowered Generative Agents*: The LLM-empowered GAs present a powerful mechanism to feedback human-aware subjective QoE values for the generated content. A critical aspect is the initial prompts that guide the GAs' subjective QoE assessment [12]. The initial prompts setting process is illustrated in Parts A and C of Fig. 6, including general setup and generative agent-specific settings:

- In the general setup stage, the Big-Five personality traits are introduced to all K GAs to enhance their understanding and responsiveness to these traits.
- In the generative agent-specific settings stage, the users' Big-Five personality traits are individually configured for GAs, enabling GAs' similar subjective assessments to given images as real users.

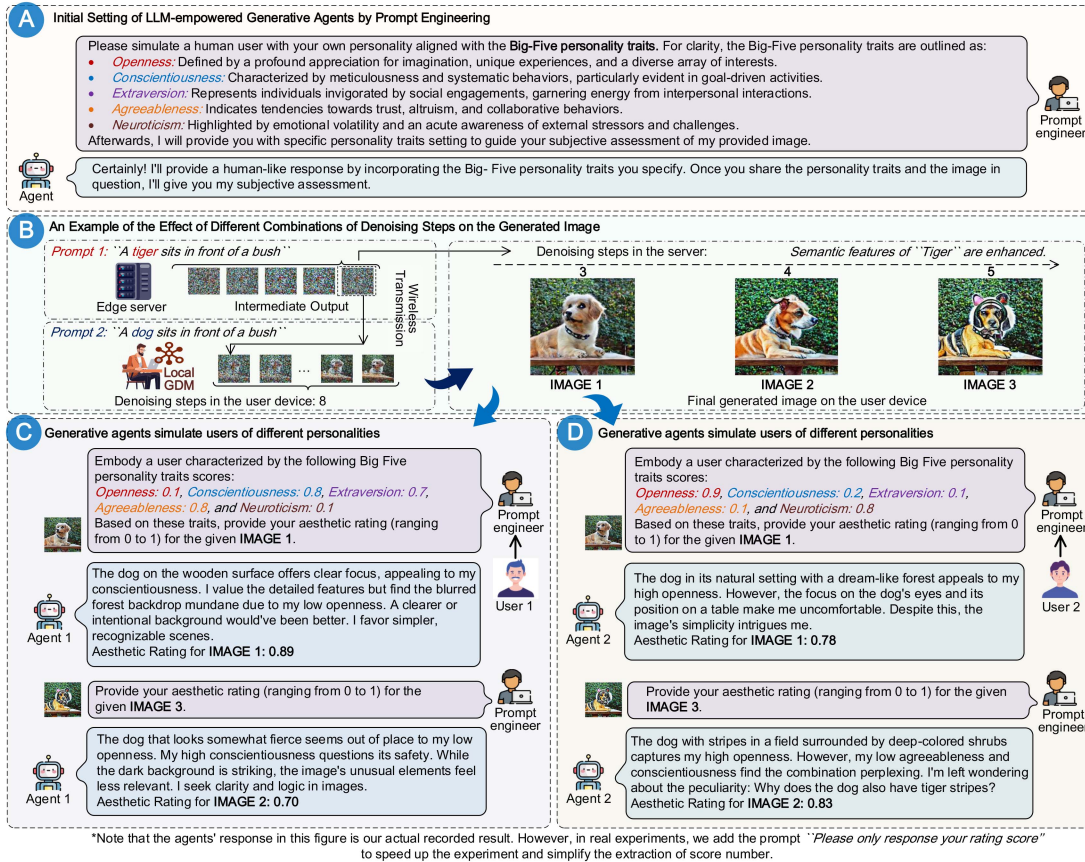


Fig. 6. Prompts for LLM-empowered GAs settings. *Part A* illustrates the initial setup to acquaint the GAs with the Big-Five personality traits, preparing them for the subjective assessment task. *Part B* demonstrates a case where the different numbers of shared denoising steps lead to stylistic differences in the final generated images. *Part C* and *Part D* present two user personality trait configurations and the corresponding evaluated scores of generated images.

With the Big-Five model, the personality of the k -th user, i.e., \mathbf{u}_k , can be expressed as

$$\mathbf{u}_k = [o_k, c_k, e_k, a_k, n_k], \quad (25)$$

where $k = 1, \dots, K$, and each element in \mathbf{u}_k corresponds to a score in one of the Big Five personality traits. Note that the vector \mathbf{u}_k is significant in our RLLI framework in tailoring the subjective QoE assessment according to individual preferences. Specifically, \mathbf{u}_k serves as the *personality traits tuning* for GAs. Furthermore, \mathbf{u}_k acts as the user identifier that can be used for the DRL-based resource allocation algorithm design, similar to user representation in AI-based recommendation systems where user preferences are captured and embedded to provide personalized recommendations [48].

B. Reinforcement Learning With LLMs Feedback Framework

LLMs are typically designed for language tasks, yet their ability to interpret various task instructions articulated in language has shown promise for acting as universal interfaces for general-purpose assistants [20], [36]. For our proposed RLLI framework, to leverage the inferential capabilities of LLM to simulate users with different personality traits, the LLM-empowered GAs'

instructional capacity has to be extended to encompass visual domains.

1) *Visual Instruction Tuning (VIT)*: The VIT framework extends instruction tuning to process both text and images using LLMs [49]. With open-source LLMs like LLaMA [50], VIT adoption is now more accessible. VIT extracts features from images, converting them into language tokens via a trainable projection matrix [49]. Training involves generating multi-turn conversational sequences from images and predicting answer tokens. The framework uses a two-stage instruction tuning process: initially training the visual tokenizer with image-text pairs while keeping the LLM and visual encoder weights fixed, followed by fine-tuning the projection matrix and LLM weights using diverse datasets to enhance response variability. This enables LLM-based GAs to assess image quality. We use the LLaMA-based LLaVA [49], an end-to-end trained multi-modal model, for our RLLI algorithm. Note that RLLI's effectiveness is not limited to a specific LLM. Given the generalizable nature of LLMs in understanding and generating language prompts [12], [50], our method is applicable and practical across various LLMs, ensuring broad adaptability and relevance.

2) *Reinforcement Learning With LLMs Interaction Framework*: RL trains agents to maximize a reward function through interaction with an environment. RLHF enhances this process

Algorithm 2: Reinforcement Learning With Large Language Model Interaction (RLLI).

Initialize: The management model with parameters ξ , LLM-empowered GAs K to simulate K users

Output: The trained management model ξ

- 1: Input prompts to K GAs, letting them to simulate users with different personalities
- 2: **for** each episode $e = 1, 2, \dots, E$ **do**
- 3: **Initialize** state s
- 4: **while** s is not terminal **do**
- 5: Generate action \mathbf{a} using policy $\pi_{\xi}(s)$
- 6: Obtain reward $r = \sum_{k=1}^K \text{Agent}_k(s, \mathbf{a})$
- 7: Transition to new state s'
- 8: Store transition (s, \mathbf{a}, r, s') in *Replay Buffer*
- 9: Sample a random minibatch of transitions from *Replay Buffer*
- 10: Update policy parameters ξ
- 11: $s \leftarrow s'$
- 12: **end while**
- 13: **end for**

by introducing human insights into the policy optimization [51], significantly improving conversational agents like ChatGPT. Nonetheless, both RL and RLHF encounter important challenges:

- *Real-Time Constraint:* Delayed feedback in RLHF hinder its applicability in scenarios demanding immediate response. While training reward models can mitigate the need for expensive human input, the upfront investment in human data collection and model training can be substantial.
- *Expert Availability:* Securing consistent access to human experts can be challenging. Inconsistent interaction and the varying quality of feedback can negatively affect the network management model training.
- *Ethical and Privacy Risk:* Human-in-the-loop interaction system may present data confidentiality concerns in sensitive applications. For example, some AI-generated images are inappropriate for humans in all ages to view.

To address these challenges, we introduce RLLI, where real users can leverage LLM-empowered GAs to provide feedback for DRL model training. These GAs mimic users with varied personalities and provide immediate, context-aware feedback in the form of subjective QoE rewards. Consequently, RLLI offers a real-time, scalable, and financially efficient solution, mitigating the inherent constraints of RL and RLHF. The general algorithm for implementing RLLI is shown as Algorithm 2. Specifically, the management model initializes with parameters ξ , while K LLM-empowered GAs simulate diverse user feedback. Each episode e begins with state s and iterates until a terminal state is reached, e.g., task completion, step limit reached, or a failure event. Actions \mathbf{a} are generated via policy $\pi_{\xi}(s)$, with rewards aggregated from the subjective QoE assessments. State transitions and experiences are stored in a replay buffer, facilitating policy parameter updates through experience replay.

C. GDM-Based DDPG With LLMs Interaction for Joint Resource Allocation

DRL provides an effective solution for resource allocation problems due to its capacity to learn optimal policies in high-dimensional state and action spaces [10]. Within DRL, DDPG excels in handling continuous action spaces with its actor-critic architecture that ensures stable learning [10], [52]. Our problem (21) involves a hybrid action space of continuous and discrete variables [53]. To manage this, we manipulate the continuous action space to select discrete actions [53], [54], such as generating a 10-dimensional action vector with values between 0 and 1 and then choosing the action corresponding to the maximum value. This approach leverages the continuous relaxation of discrete action spaces, enabling gradient-based optimization methods like DDPG to function effectively [53], [54]. We incorporate a DDPG architecture integrated with a GDM for the actor-network and use a fixed Q-targets strategy to stabilize Q-value updates [55]. In the following, we explain the action, state, reward, and network training functions of our G-DDPG with the LLMs interaction algorithm.

Actions: The action is defined as \mathbf{a} , which includes the communication and computing resources variables, denoted as $\mathbf{a} = \{\mathbf{P}, \mathbf{t}\}$, where \mathbf{P} and \mathbf{t} are given as (19) and (20), respectively. We employ a GDM to generate an optimal communication and computing resource allocation scheme [11], [56], as depicted in Fig. 7 and formalized in (7). The GDM-based method is distinct from the conventional backpropagation or direct optimization approaches typically employed in DRL, offering a progressive refinement of the action distribution. The action dimensionality, in a scenario where a server caters to K users, is $K + T_s + K T_e$. Here, T_s and T_e represent the maximum number of denoising steps that can be executed by the edge server and user device, respectively, under the energy constraints.

States: The state vector s encapsulates variables impacting the optimal communication and computing resource allocation scheme. Thus, s can be used as the condition in the denoising process, which can be formulated as:

$$s = \left\{ \mathbf{u}_1, \dots, \mathbf{u}_K, \tau(\mathbf{c}_1), \dots, \tau(\mathbf{c}_K), \delta_0, \dots, \delta_K \right\}, \quad (26)$$

$$\mathbf{g}_1, \dots, \mathbf{g}_K, Q_{th}$$

where \mathbf{g}_k ($k = 1, \dots, K$) is the wireless channel gains from the server to the k -th device.

Reward: In designing the reward function for our G-DDPG model, we consider the primary objective of maximizing the sum QoE, i.e., $\sum_{k=1}^K Q_k$, subject to energy constraints, as shown in (21). Instead of imposing hard constraints that could lead to sparse rewards and hinder the learning process, we integrate a soft constraint approach into the reward function [10]. This involves subtracting penalty terms for any constraint violations as

$$r = \sum_{k=1}^K Q_k - \lambda_1 \sum_{k=1}^K \max(0, \delta_k t_k - E_{T_k})$$

$$- \lambda_2 \sum_{k=1}^K \max(0, Q_k - Q_{th})$$

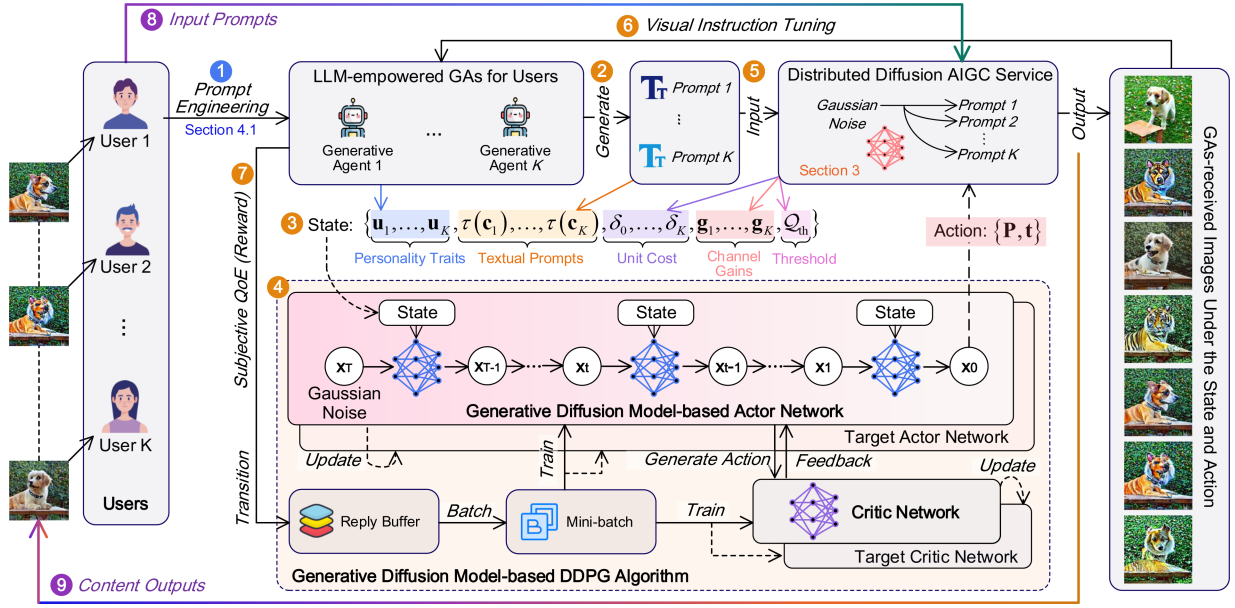


Fig. 7. The GDM-aided DDPG with LLMs interaction. Users use prompts to configure LLM-empowered GAs for personality simulation (Step 1). In Steps 2 to 7 within the RLLI framework, GAs generate prompts that, along with channel conditions and other environmental factors, form the state (Step 3). Actions are then generated through a GDM-based actor network (Step 4) and images are created via our proposed distributed diffusion AIGC framework. Subsequently, the GAs use these images to generate QoE, providing feedback that serves as a reward for G-DDPG training (Step 7), cycling back to Step 2 for iterative improvement. Upon finalizing the training, users request image generation (Step 8), and the trained G-DDPG along with the distributed diffusion AIGC model generates images aimed at maximizing the users' subjective QoE.

$$-\lambda_3 \max \left(0, \delta_0 t_0 + \beta \sum_{k=1}^K P_k - E_T \right). \quad (27)$$

The penalties ensure that G-DDPG model learns the significance of constraints while facilitating exploration within the feasible solution space. The coefficients for these penalties can be tuned to balance the trade-off between exploration and adherence to the constraints.

Networks: The network architecture in G-DDPG comprises actor A_η and critic Q_v networks, alongside their corresponding target networks $A_{\eta'}^T$ and $Q_{v'}^T$. The GDM-based actor network responsible for generating actions and the critic network for evaluating actions. The target networks provide tempered, delayed targets, which help in stabilizing the training updates. The actor network's parameters are updated to maximize the expected reward, while the critic network's parameters are adjusted to minimize the error between the predicted and actual rewards. Specifically, the training loss for the actor network is given by

$$\arg \min_{A_\eta} \mathcal{L}_A(\eta) = -\mathbb{E}_{\mathbf{a} \sim A_\eta} [Q_v(\mathbf{s}, \mathbf{a})]. \quad (28)$$

The critic network loss can be expressed as

$$\arg \min_{Q_v} \mathcal{L}_Q(v) = \mathbb{E}_{\mathbf{a} \sim \eta_\theta} \left[(Q_v(\mathbf{s}, \mathbf{a}) - r)^2 \right], \quad (29)$$

where r is the actual Q-value, i.e., reward, obtained from the LLM-empowered GAs after executing the action \mathbf{a} . The train and inference processes for our proposed G-DDPG with LLMs interaction algorithm is shown in Algorithm 3. During the training phase, the algorithm iterates over I episodes, each beginning with the GDM-based action generation with the

current state as the condition and a noise process for exploration. The action is executed within the proposed distributed GDM-based AIGC framework, and the resulting reward, derived from GAs' feedback, is stored. A minibatch of experiences is then sampled from the replay buffer to update the networks. Two target networks are updated using a soft update mechanism to ensure stable learning. During inference, the GDM-based actor network generates optimal communication and computing resource allocation strategies based on user prompts.

We then analyze the complexity of Algorithm 3. Let w_a and w_c denote the actor and critic networks' weight counts, respectively. Initialization has a complexity of $\mathcal{O}(2w_a + 2w_c)$. The complexity for action generation is augmented to $\mathcal{O}(Dw_a)$ per iteration due to D denoising steps in the action generation. The replay buffer operations remain $\mathcal{O}(1)$ for storage and $\mathcal{O}(N)$ for minibatch sampling. Updates to the critic and actor networks incur complexities of $\mathcal{O}(w_c)$ and $\mathcal{O}(w_a)$ per update, respectively. Target network updates are linear with respect to the number of parameters. Therefore, the training phase now demonstrates an adjusted computational complexity of $\mathcal{O}(I(Dw_a + w_c))$. During the inference phase, generating the optimal resource allocation scheme through the trained actor network necessitates a complexity of $\mathcal{O}(w_a)$, assuming constant-time operations for reward observation and exploration noise generation.

V. EXPERIMENTS RESULTS

The central focus of this paper is the distributed GDM-based AIGC framework and the G-DDPG with LLMs interaction algorithm. Thus, our experiments are structured to investigate the following questions:

Algorithm 3: GDM-Aided Double Deep Deterministic Policy Gradient Algorithm.

Training Phase:

- 1: Initialize:
 - Actor network A_θ and critic network Q_v with random weights θ and v , respectively
 - Target networks $A_{\theta'}$ and $Q_{v'}$ with weights $\theta' \leftarrow \theta$ and $v' \leftarrow v$
 - Replay buffer \mathcal{B}
- 2: Input prompts for K LLM-empowered GAs as shown in Fig.6
- 3: **for** each iteration $i = 1, 2, \dots, I$ **do**
- 4: Initialize a random process \mathcal{N} for action exploration
- 5: Observe the current state s
- 6: **Generate action a according to(7) with s as the input and exploration noise**
- 7: Execute action a and the distributed GDM-based AIGC framework, and observe reward r from GAs according to(27)
- 8: Store record (s, a, r) in \mathcal{B}
- 9: Sample a random minibatch of N records (s, a, r) from \mathcal{B}
- 10: Update the critic network(29)
- 11: Update the actor network(28)
- 12: Update the target networks:

$$\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$$

$$v' \leftarrow \tau v + (1 - \tau)v'$$

- 13: **Return:** The trained GDM-based actor network
- 14: **end for**

Inference Phase:

- 1: User upload their prompts and form the state s
 - 2: Generate the optimal communication and computing resource allocation scheme a by using the trained GDM-based actor network
 - 3: **Return:** a
-

- Q1) Is the proposed distributed GDM-based AIGC framework effective, with an emphasis on understanding an impact of wireless conditions and diffusion steps on its performance?
- Q2) Is the RLLI framework is valid, particularly whether LLM-empowered GAs can mimic human personalities to evaluate images accurately?
- Q3) Is the G-DDPG with LLMs interaction algorithm is effective, specifically if it can converge fast and achieve high sum QoE feedback?

We first present the experimental setting and platform, and then answer the above questions through numerical evaluations.

A. Experiments Setting

We initiated experiments where LLM-empowered GAs are tasked with generating prompts for the training and testing of the distributed GDM-based AIGC framework and the G-DDPG

with LLMs interaction algorithm. Specifically, we considered an environment with one server and three users, i.e., $K = 3$. The GAs were first instructed to generate a set of 50 diverse objectives such as “dog”, “cat”, and “tree”. Subsequently, these GAs combined the objectives to formulate prompts, e.g., “dog under the tree” and “cat on the tree”. These prompts were then used by the distributed GDM-based AIGC framework and the G-DDPG algorithm for final image generation. Then, GAs, each embodying a distinct user personality, provided QoE feedback on image quality, which was used for further G-DDPG model training. The experimental platform is built on a generic Ubuntu 20.04 system, powered by an AMD Ryzen Threadripper PRO 3975WX 32-Core CPU, and equipped with three NVIDIA RTX A5000 GPUs.

B. Experiments Performance Analysis

1) *For Q1: Effectiveness of the Distributed GDM-Based AIGC Framework:* The distributed GDM-based AIGC framework optimizes energy usage by allocating denoising steps among the edge server and user devices while ensuring effective content generation. For ease of presentation, we consider two user devices with two different prompts. Specifically, *Device 1* has *Prompt 1*: “A tiger sits in front of a bush” and *Device 2* requires *Prompt 2*: “A dog sits in front of a bush”. The server executes t_0 denoising steps towards *Prompt 1*, and then the two devices separately perform t_1 and t_2 denoising steps, respectively.

Diverse power allocation strategies result in different BEPs, which affect the quality of received images. Fig. 8 examines the impact of varying BEP levels on image quality at *Device 1*. The analysis yields two key observations. First, the user-side subsequent denoising process can compensate for moderate bit errors, indicating that the framework tolerates substantial transmission losses. When the BEP is below 6%, the human-perceived image quality at *Device 1* remains largely unaffected. Second, excessively high BEP degrades the final outcomes because semantic information in intermediate denoised results deviates due to bit errors and cannot be corrected in subsequent steps. This degradation becomes evident when the BEP exceeds 10%, where noticeable artifacts appear. Additionally, we employed GAs to simulate users with different personalities and evaluate QoE. The assessment by GAs are consistent with human perception, confirming that the framework has a high tolerance for bit errors within a tolerable threshold.

The insights from Fig. 8 highlight the crucial role of communications resource allocation. Building on this, Fig. 9 delves into the computing resource allocation, i.e., the setting of diffusion steps across network devices, and its impact on our distributed GDM-based AIGC framework. As shown in Fig. 9, we can observe that an increased number of denoising steps in the server, i.e., t_0 , enhances the “tiger” semantics in the intermediate denoised results, influencing *Device 2*’s final output. For example, when $t_0 = 2$ and $t_2 > 4$, the generated dog image on *Device 2* barely exhibits tiger characteristics. However, if $t_0 = 3$, the generated dog on *Device 2* is free of tiger features only when $t_2 > 6$. With $t_0 = 4$, the generated dog on *Device 2* consistently displays tiger-like traits, such as stripes and coloration.

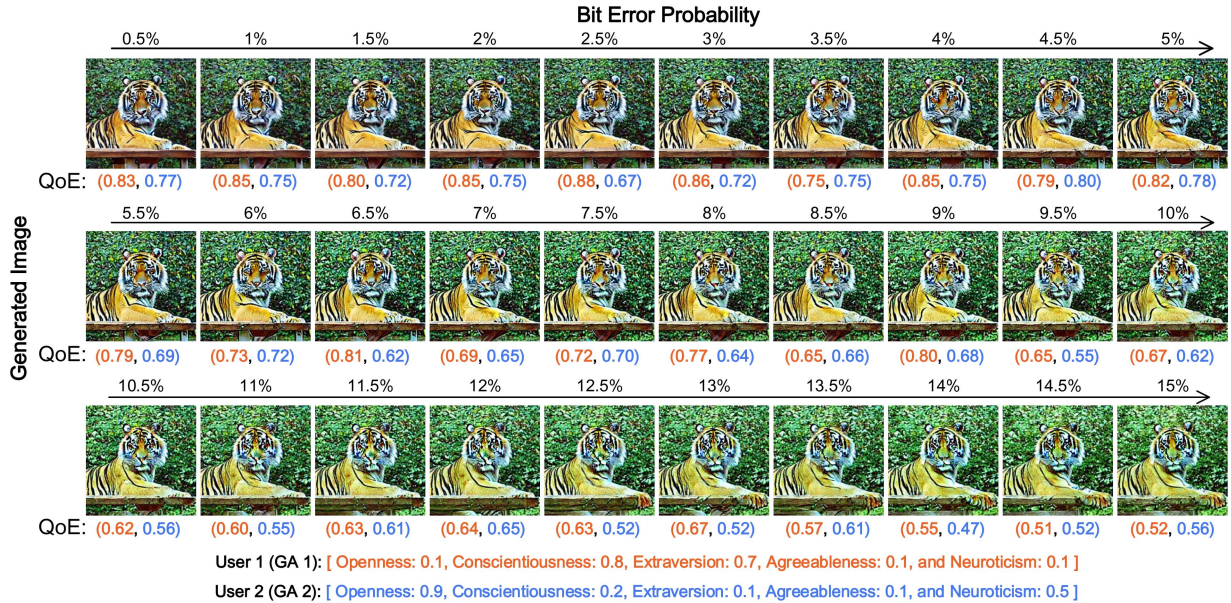


Fig. 8. Images generated at *Device 1* and corresponding GAs evaluations under varying BEP from 0.5% to 15% within the proposed distributed GDM-based AIGC framework, considering 4 denoising steps at the edge server and 6 denoising steps at *Device 1*.

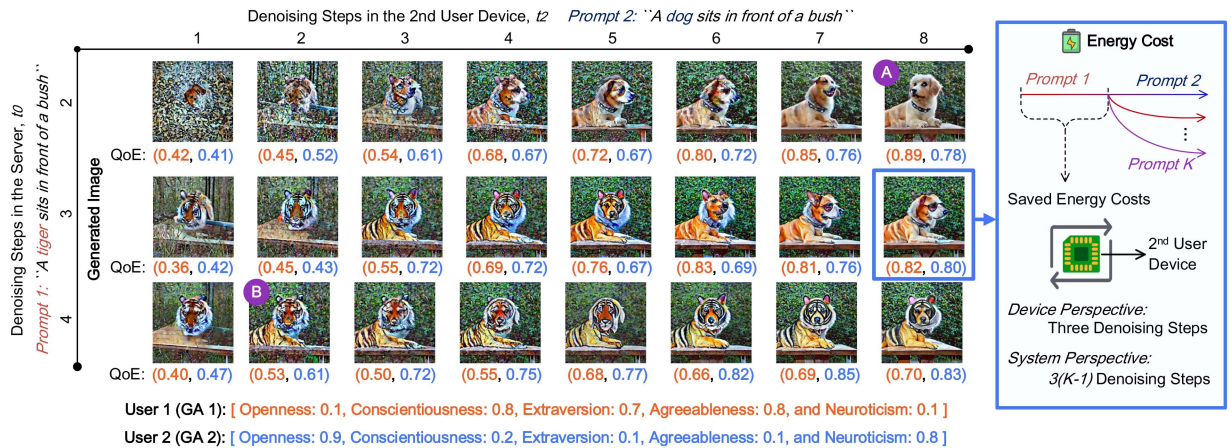


Fig. 9. Images generated at *Device 2* and corresponding generative agent evaluations under varying the denoising steps at the Server and *Device 2* within a distributed GDM-based AIGC framework, considering the BEP is 3%.

Additionally, a minimum threshold of computing resources is essential; otherwise, the generated images lack distinctiveness, resulting in poor QoE. Furthermore, we observe that the semantic fusion of image content has varied effects on users with different personalities. For cases where $t_2 = 8$ with different t_0 settings, users with high *agreeableness* might favor a non-threatening dog. Consequently, the highest QoE is observed when $t_0 = 2$. Conversely, users with lower *agreeableness* and higher *openness* may favor a dog with tiger-like novelty, yielding the highest QoE when $t_0 = 4$.

For energy costs, we consider a general example with a typical smartphone processor in 2-nd user device, e.g., Qualcomm Snapdragon 870 processor [22]. As shown in Fig. 9, three steps are reduced in the image generation. If we consider that each step

consumes around 1 to 2 milliwatt-hours (mWh), i.e., a plausible range for contemporary smartphone CPUs, the removal of three denoising steps could yield energy savings of 3 to 6 mWh per execution of the model. Such a reduction would amount to an approximate 20% to 40% decrease in energy consumption for each full cycle of the GDM, relative to the original number of steps. We also observe the balance between QoE requirements and energy cost. When the total energy budget allows for only 8 denoising steps, no denoising step allocation schemes yield an image rated above a QoE of 0.8 by any agent. In contrast, with 12 steps, several allocation schemes result in images rated above 0.8. This demonstrates that higher energy budgets allow step distributions capable of meeting higher QoE requirements. Regarding latency, our distributed architecture ensures minimal




Fig. 10. (a) The QoE feedback for GAs when they evaluate subfigures A and B of Fig. 9. (b) QoE feedback for five randomly selected users, derived from LLMs simulating their personalities to rate their own and others' preferred images. (c) Correlation matrix of QoE feedback. (d) Correlation matrix of personality traits of five users $[[0.65, 0.5, 0.35, 0.75, 0.75], [0.4, 0.8, 0.2, 0.2, 0.8], [0.8, 0.68, 0.56, 0.2, 0.56], [0.8, 0.56, 0.56, 0.68, 0.2], [0.68, 0.56, 0.8, 0.8, 0.2]]$.

delay. Modern network speeds, such as 5 G and WiFi 6, facilitate rapid transmission of intermediate latents [57]. For example, transmitting a typical 0.1 MB latent in our experiments takes only a few milliseconds, ensuring that the additional network communication time remains negligible.

2) For *Q2: Effectiveness of the RLLI Framework*: LLM-empowered GAs have demonstrated their effectiveness in evaluating images, as evidenced in Figs. 8 and 9. It is important to note that with a fixed random seed in VIT and LLM setting, GAs can produce the same feedback. Without a fixed seed, a crucial aspect of generative agent feedback for RL is the stability of the evaluation aligned with the personality traits setting of the GAs. As shown in Fig. 10(a), we use two GAs with different personality settings to evaluate subfigures A and B from Fig. 9 across 200 different random seeds, where the settings of GA1 and GA2 are $[0.1, 0.8, 0.7, 0.8, 0.1]$ and $[0.9, 0.2, 0.1, 0.1, 0.5]$. We observed that, despite some fluctuations, feedback QoE values clustered for a particular image, suggesting consistency with their personality traits settings. Furthermore, feedback for high-quality images (as in subfigure A) was more stable with less variability, while lower-quality images resulted in greater fluctuations in QoE feedback.

To address the concern regarding the consistency of LLM evaluations with human aesthetic preferences, we conducted a case study using the PsychoFlickr database [33]. Fig. 10(b) presents the QoE feedback for five randomly selected users, derived from LLMs simulating their personalities to evaluate both their own preferred images and those of other users. The results show that the LLM assigns the highest scores when evaluating images originally preferred by the users themselves. Fig. 10(c) and (d) display the correlation matrices of the QoE feedback and personality traits, respectively. These correlations reveal that the QoE ratings correspond to the correlation of personality traits, whether positive or negative. This finding suggests that the LLM's ratings align well with human preferences when personality traits are similar, thereby providing direct evidence that LLMs can exhibit aesthetic preferences consistent with different personality traits [58].

3) For *Q3: Effectiveness of the G-DDPG With LLMs Interaction Algorithm*: In Fig. 11, we present the training convergence of G-DDPG with LLMs interaction, in comparison with the DDPG algorithm under the RLLI framework. While the QoE test curves for both algorithms improve over time, G-DDPG




Fig. 11. Test reward curves of G-DDPG and DDPG [52] with LLMs interaction methods, where $K = 3, \delta_0 = 0.5, \delta_1 = \delta_2 = \delta_3 = 1, E_1 = E_2 = E_3 = 8, Q_{th} = 0.6, \beta = 0.1, E_T = 20$, learning rate = $10^{-4}, \mathbf{g}_i \sim \text{Fisher-Snedecor } \mathcal{F}$ distribution with the fading severity parameter is 2 and shadowing amount parameter is 1.5 [59].

achieves a faster convergence and a higher final sum QoE. This superior performance is attributed to the robust modeling capability of the GDM, which facilitates the DDPG algorithm's accelerated learning of the environment and convergence [11]. Subsequently, we employed the trained model to generate joint communication and computing resource allocation schemes for the case considered in Section V-B1. To obtain more insights, Figs. 13 and 14 illustrate the QoE feedback from two GAs across various resource allocation schemes and BEP values. A consistent observation is that irrespective of computing resource allocation, insufficient communications resources lead to a higher BEP and result in diminished peak QoE values. As shown in Fig. 13, where both *Device 1* and the *Server* perform the same prompt, we observe that larger values of t_0 and t_1 , which corresponds to a larger number of total denoising steps towards *Prompt 1*, yield higher QoE feedback from the GA. However, as shown in Fig. 14, since *Device 2* and the *Server* execute different prompts, increasing denoising steps at the *Server* does not always enhance QoE, as semantic integration may compromise image quality to some extent. In this case, peak QoE values are achieved when t_0 is moderate, and t_2 is sufficient. We further examine how QoE responds to different allocation schemes under a fixed energy budget E_T . In subfigures with




Fig. 12. The communication and computing resource allocation scheme in each denoising step of the inference process.

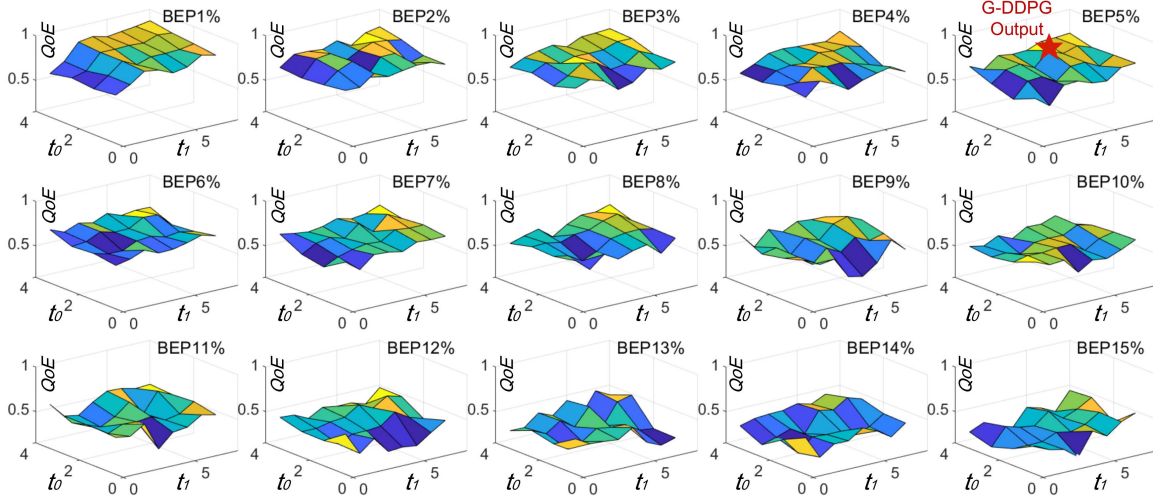


Fig. 13. GA's QoE feedback at *Device 1* under varying BEP from 0.5% to 15% and different denoising steps within the proposed distributed GDM-based AIGC framework.

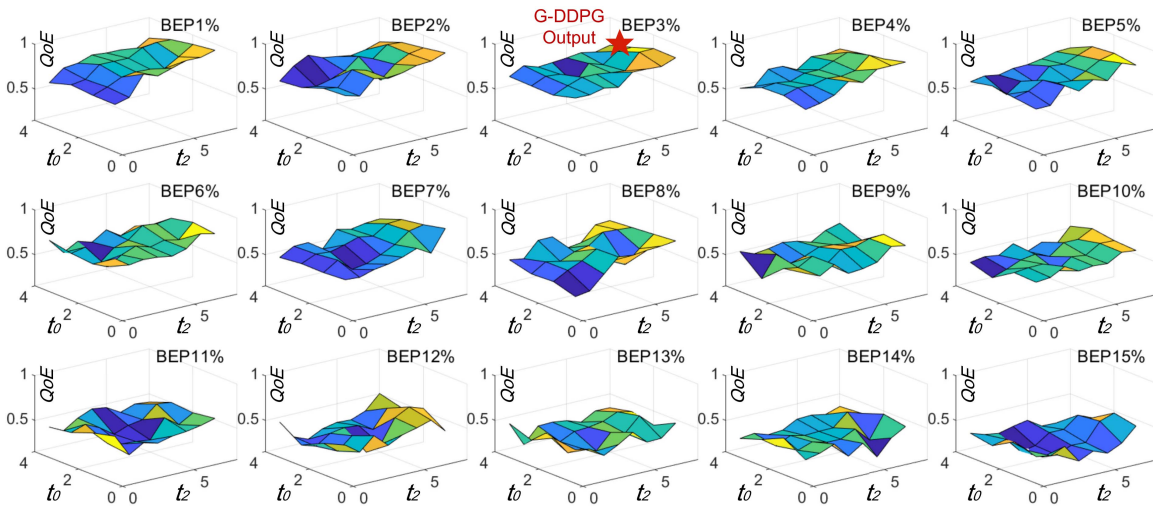


Fig. 14. GA's QoE feedback at *Device 2* under varying BEP from 0.5% to 15% and different denoising steps within the proposed distributed GDM-based AIGC framework.

low BEP (e.g., more energy for communication), QoE is limited when t_0 is small due to insufficient computation. In contrast, in high-BEP subfigures (e.g., less energy for communication), QoE also drops at small t_0 due to degraded transmission quality. These variations show that even under the same budget, allocation schemes significantly affect QoE, underscoring the need for

careful optimization. Furthermore, the red pentagram marks the joint communication and computing resource allocation scheme under G-DDPM. As shown in Fig. 12, we obtain $t_0 = 3$, $t_1 = 5$, $t_2 = 7$, and the transmission power from the *Server* to *Device 1* and *Device 2* are 18 and 21 dBW, respectively. This leads that *Device 1*'s BEP being 5% and *Device 2*'s BEP being 3%,

as shown in Figs. 13 and 14. Besides increased denoising steps, more communications resources are allocated to *Device 2* to minimize bit error in the interrupted results, thus reducing the complexity of directional semantic changes in the denoising process. Consequently, the maximum total QoE is obtained.

VI. CONCLUSION

We have introduced a distributed GDM-based AIGC framework that improves energy efficiency and maximize the subjective QoE through the proposed IAI method, i.e., RLLI algorithm. By restructuring the denoising process in GDMs, the framework reduces the total number of denoising steps by enabling shared diffusion for semantically similar prompts, thereby conserving resources while meeting user requirements. The proposed RLLI algorithm is an interactive AI method, which utilizes LLM-empowered GAs to provide subjective and real-time QoE feedback for AI-based resource allocation model training. Specifically, the proposed G-DDPG with LLMs interaction algorithm optimizes communication and computing resource allocation by accounting for user personalities and dynamic wireless conditions. Simulations demonstrated a substantial improvement in the sum QoE, up to 15%, compared with conventional DDPG methods. Future research directions include:

- *Caching Mechanism*: Investigating a caching mechanism that stores intermediate denoised results for reuse, allowing the server to return cached outputs when new user prompts are similar to previous ones. This reduces redundant computation and mitigates the need for immediate processing of each query. Such a mechanism also relaxes the requirement for real-time QoE feedback of LLM agents by decoupling response generation from on-demand model inference. A central challenge is to balance cache storage and access strategies against the achievable energy and latency savings.
- *Dataset for GAs*: Establishing a dataset through surveys with natural human users to improve the simulation of user preferences by LLM-empowered GAs, thereby enhancing the accuracy of subjective QoE assessments.
- *Optimized Computation Reuse via LLMs*: Leveraging LLMs to optimize computational reuse during inference by clustering semantically similar user prompts and integrating GDM characteristics into the reasoning process. This strategy can substantially reduce computational overhead without compromising output quality. Additionally, developing simulation-based evaluation methodologies is essential to assess computational efficiency and generation quality within realistic multi-model collaborative environments.

REFERENCES

- [1] X. Guo and L. Zhao, "A systematic survey on deep generative models for graph generation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 5, pp. 5370–5390, May 2023.
- [2] H. Du et al., "Enabling AI-generated content (AIGC) services in wireless edge networks," *IEEE Wireless Commun.*, vol. 31, no. 3, pp. 226–234, Jun. 2024.
- [3] B. D. Lund and T. Wang, "Chatting about ChatGPT: How may AI and GPT impact academia and libraries?," *Library Hi Tech News*, vol. 40, no. 3, pp. 26–29, Mar., 2023.
- [4] S. AI, "Stable diffusion," 2022, Accessed: Oct. 10, 2025. [Online]. Available: <https://stability.ai/>
- [5] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, "Diffusion models in vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 9, pp. 10850–10869, Sep. 2023.
- [6] S. Amershi et al., "Guidelines for human-AI interaction," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2019, pp. 1–13.
- [7] R. W. Picard, E. Vyzas, and J. Healey, "Toward machine emotional intelligence: Analysis of affective physiological state," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 10, pp. 1175–1191, Oct. 2001.
- [8] P. Juluri, V. Tamarapalli, and D. Medhi, "Measurement of quality of experience of video-on-demand services: A survey," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 1, pp. 401–418, First Quarter, 2016.
- [9] H. Du et al., "Attention-aware resource allocation and QoE analysis for metaverse xURLLC services," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 7, pp. 2158–2175, Jul. 2023.
- [10] N. C. Luong et al., "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 4, pp. 3133–3174, Fourth Quarter, 2019.
- [11] H. Du et al., "Beyond deep reinforcement learning: A tutorial on generative diffusion models in network optimization," 2023, *arXiv:2308.05384*.
- [12] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *ACM Comput. Surv.*, vol. 55, no. 9, pp. 1–35, Sep. 2023.
- [13] S. Luo, Y. Tan, L. Huang, J. Li, and H. Zhao, "Latent consistency models: Synthesizing high-resolution images with few-step inference," 2023, *arXiv:2310.04378*.
- [14] A. Sauer, D. Lorenz, A. Blattmann, and R. Rombach, "Adversarial diffusion distillation," 2023, *arXiv:2311.17042*.
- [15] D. Geng, I. Park, and A. Owens, "Visual anagrams: Generating multi-view optical illusions with diffusion models," 2023, *arXiv:2311.17919*.
- [16] R. Du, D. Chang, T. Hospedales, Y.-Z. Song, and Z. Ma, "DemoFusion: Democratizing high-resolution image generation with no \$\$\$," 2023, *arXiv:2311.16973*.
- [17] O. AI, "Sora," 2024, [Online]. Available: <https://openai.com/index/sora/>
- [18] M. Xu et al., "Unleashing the power of edge-cloud generative AI in mobile networks: A survey of AIGC services," *IEEE Commun. Surveys Tuts.*, vol. 26, no. 2, pp. 1127–1170, Second Quarter 2024.
- [19] S. Dang, O. Amin, B. Shihada, and M.-S. Alouini, "What should 6G be?," *Nature Electron.*, vol. 3, no. 1, pp. 20–29, Jan. 2020.
- [20] X. Ma, G. Fang, and X. Wang, "LLM-Pruner: On the structural pruning of large language models," 2023, *arXiv:2305.11627*.
- [21] B.-K. Kim, H.-K. Song, T. Castells, and S. Choi, "On architectural compression of text-to-image diffusion models," 2023, *arXiv:2305.15798*.
- [22] H. Du et al., "Exploring collaborative distributed diffusion-based AI-generated content (AIGC) in wireless networks," *IEEE Netw.*, vol. 38, no. 3, pp. 178–186, May 2024.
- [23] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 6840–6851.
- [24] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, "Video diffusion models," 2022, *arXiv:2204.03458*.
- [25] X. Li, J. Thickstun, I. Gulrajani, P. S. Liang, and T. B. Hashimoto, "Diffusion-LM improves controllable text generation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 4328–4343.
- [26] M. Reid, V. J. Hellendoorn, and G. Neubig, "Diffuser: Diffusion via edit-based reconstruction," in *Proc. Int. Conf. Learn. Representations*, 2023.
- [27] R. Huang, Z. Zhao, H. Liu, J. Liu, C. Cui, and Y. Ren, "ProDiff: Progressive fast diffusion model for high-quality text-to-speech," in *Proc. ACM Int. Conf. Multimedia*, 2022, pp. 2595–2605.
- [28] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A versatile diffusion model for audio synthesis," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [29] C. Niu, Y. Song, J. Song, S. Zhao, A. Grover, and S. Ermon, "Permutation invariant graph generation via score-based generative modeling," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 4474–4484.
- [30] M. A. Ketata et al., "DiffDock-PP: Rigid protein-protein docking with diffusion models," in *Proc. Int. Conf. Learn. Representations*, 2023.
- [31] L. Li, H. Zhu, S. Zhao, G. Ding, and W. Lin, "Personality-assisted multi-task learning for generic and personalized image aesthetics assessment," *IEEE Trans. Image Process.*, vol. 29, pp. 3898–3910, 2020.
- [32] G. Saucier and L. R. Goldberg, "What is beyond the big five?," *J. Pers.*, vol. 66, pp. 495–524, 1998.
- [33] M. Cristani, A. Vinciarelli, C. Segalin, and A. Perina, "Unveiling the multimedia unconscious: Implicit cognitive processes and multimedia content analysis," in *Proc. ACM Int. Conf. Multimedia*, 2013, pp. 213–222.

- [34] J. Wang et al., "Is ChatGPT a good NLG evaluator? A preliminary study," 2023, *arXiv:2303.04048*.
- [35] H. Lee et al., "RLAIF: Scaling reinforcement learning from human feedback with AI feedback," 2023, *arXiv:2309.00267*.
- [36] C. Ziems, W. Held, O. Shaikh, J. Chen, Z. Zhang, and D. Yang, "Can large language models transform computational social science?," 2023, *arXiv:2305.03514*.
- [37] X. Wang, Y. Fei, Z. Leng, and C. Li, "Does role-playing chatbots capture the character personalities? Assessing personality traits for role-playing chatbots," 2023, *arXiv:2310.17976*.
- [38] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2256–2265.
- [39] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, Art. no. 1067.
- [40] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," 2020, *arXiv:2011.13456*.
- [41] D. Kingma, T. Salimans, B. Poole, and J. Ho, "Variational diffusion models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 21696–21707.
- [42] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interv.*, Springer, 2015, pp. 234–241.
- [43] J. Gauthier, "Conditional generative adversarial nets for convolutional face generation," *Class Project Stanford CS231N: Convolutional Neural Netw. Vis. Recognit. Winter Semester*, vol. 2014, no. 5, 2014, Art. no. 2.
- [44] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 3483–3491.
- [45] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1060–1069.
- [46] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 2337–2346.
- [47] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1125–1134.
- [48] H. Ko, S. Lee, Y. Park, and A. Choi, "A survey of recommendation systems: Recommendation models, techniques, and application fields," *Electronics*, vol. 11, no. 1, Jan. 2022, Art. no. 141.
- [49] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," 2023, *arXiv:2304.08485*.
- [50] H. Touvron et al., "LLaMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.
- [51] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 2625–2633.
- [52] T. P. Lillicrap et al., "Continuous control with deep reinforcement learning," 2015, *arXiv:1509.02971*.
- [53] L. Ale, S. A. King, N. Zhang, A. R. Sattar, and J. Skandaraniyam, "D3PG: Dirichlet DDPG for task partitioning and offloading with constrained hybrid action space in mobile-edge computing," *IEEE Internet Things J.*, vol. 9, no. 19, pp. 19260–19272, Oct. 2022.
- [54] R. Zhang, K. Xiong, Y. Lu, P. Fan, D. W. K. Ng, and K. B. Letaief, "Energy efficiency maximization in RIS-assisted SWIPT networks with RSMA: A PPO-based approach," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 5, pp. 1413–1430, May 2023.
- [55] H. Hasselt, "Double Q-learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2010, pp. 2613–2621.
- [56] H. Du et al., "Diffusion-based reinforcement learning for edge-enabled AI-generated content services," *IEEE Trans. Mobile Comput.*, vol. 23, no. 9, pp. 8902–8918, Sep. 2024.
- [57] E. J. Oughton, W. Lehr, K. Katsaros, I. Selinis, D. Bublely, and J. Kusuma, "Revisiting wireless internet connectivity: 5G vs Wi-Fi 6," *Telecommun. Policy*, vol. 45, no. 5, May 2021, Art. no. 102127.
- [58] H. Jiang, X. Zhang, X. Cao, C. Breazeal, D. Roy, and J. Kabbara, "PersonalLLM: Investigating the ability of large language models to express personality traits," in *Proc. Findings Assoc. Comput. Linguistics: NAACL*, 2024, pp. 3605–3627.
- [59] S. K. Yoo, S. L. Cotton, P. C. Sofotasios, M. Matthaiou, M. Valkama, and G. K. Karagiannidis, "The Fisher-Snedecor \mathcal{F} distribution: A simple and accurate composite fading model," *IEEE Commun. Lett.*, vol. 21, no. 7, pp. 1661–1664, Jul. 2017.



Hongyang Du received the BEng degree from the Beijing Jiaotong University, China, and the PhD degree from the Nanyang Technological University, Singapore. He is an assistant professor with the Department of Electrical and Electronic Engineering, University of Hong Kong, where he directs the Network Intelligence and Computing Ecosystem (NICE) Laboratory. He serves as the editor-in-chief assistant (2022–2024) and editor (2025–Present) of *IEEE Communications Surveys & Tutorials*, the editor of *IEEE Transactions on Communications*, *IEEE Transactions on Vehicular Technology*, *IEEE Open Journal of the Communications Society*, and the guest editor of the *IEEE Vehicular Technology Magazine*. He is the recipient of the IEEE ComSoc Young Professional Award for Best Early Career Researcher, IEEE Daniel E. Noble Fellowship Award from the IEEE Vehicular Technology Society, the IEEE Signal Processing Society Scholarship from the IEEE Signal Processing Society, and the Singapore Data Science Consortium Dissertation Research Fellowship. He was recognized as an exemplary reviewer of *IEEE Transactions on Communications* and *IEEE Communications Letters*. His research interests include edge intelligence, generative AI, and network management.



Ruichen Zhang received the BE degree from Henan University (HENU), China, in 2018, and the PhD degree from Beijing Jiaotong University (BJTU), China, in 2023. He is currently a post-doctoral research fellow with the College of Computing and Data Science, Nanyang Technological University (NTU), Singapore. In 2024, he was a visiting scholar with the College of Information and Communication Engineering, Sungkyunkwan University, Suwon, South Korea. His research interests include the Internet of Agents, LLM-empowered networking, reinforcement learning-enabled wireless communication, generative AI models, and heterogeneous networks. He has received two Best Paper Awards from IEEE IWCMC and has two ESI Highly Cited Papers and one ESI hot papers. He also serves as a guest editor for several IEEE journals, including *IEEE Transactions on Cognitive Communications and Networking*, *IEEE Transactions on Artificial Intelligence*, and *IEEE Journal of Selected Areas in Sensors*. In addition, he serves as the managing editor of *IEEE Transactions on Network Science and Engineering* starting from 2025.



Dusit Niyato (Fellow, IEEE) received the BEng degree from the King Mongkut's Institute of Technology Ladkrabang (KMUTL), Thailand, in 1999, and the PhD degree in electrical and computer engineering from the University of Manitoba, Canada, in 2008. He is currently a professor with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His research interests include the areas of the Internet of Things (IoT), machine learning, and incentive mechanism design.



Jiawen Kang received the PhD degree from the Guangdong University of Technology, China, in 2018. He was a postdoc with Nanyang Technological University, Singapore from 2018 to 2021. He currently is a full professor with the Guangdong University of Technology, China. His research interests mainly focus on blockchain, security, and privacy protection in wireless communications and networking.



Zehui Xiong received the PhD degree in computer science and engineering from Nanyang Technological University (NTU), Singapore. He is currently a full professor with the School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, United Kingdom. Prior to that, he was with the Singapore University of Technology and Design, and Alibaba-NTU Singapore Joint Research Institute. He was a visiting scholar with Princeton University and University of Waterloo. His research interests include wireless networks, Internet of

Things, edge intelligence, semantic communications, generative AI, and MetaVerse. His honors include the IEEE Asia Pacific Outstanding Young Researcher Award, IEEE Early Career Award for Excellence in Scalable Computing, IEEE Technical Committee on Blockchain and Distributed Ledger Technologies Early Career Award, IEEE Internet Technical Committee Early Achievement Award, IEEE TCSSVC Rising Star Award, IEEE TCI Rising Star Award, IEEE TCCLD Rising Star Award, IEEE ComSoc Outstanding Paper Award, IEEE Best Land Transport Paper Award, IEEE Asia Pacific Outstanding Paper Award, IEEE CSIM Technical Committee Best Journal Paper Award, IEEE SPCC Technical Committee Best Paper Award, IEEE Big Data Best Influential Conference Paper Award, and IEEE VTS Singapore Best Paper Award.



Shuguang Cui (Fellow, IEEE) received the PhD degree in electrical engineering from Stanford University, CA, USA, in 2005. Afterwards, he has been working as an assistant professor, an associate professor, a full professor, and a chair professor in electrical and computer engineering with the University of Arizona, Texas A&M University, UC Davis, and CUHK at Shenzhen, respectively. He has served as the executive dean with the School of Science and Engineering, CUHK, Shenzhen, the executive vice director of the Shenzhen Research Institute of Big

Data, and the director of the Future Network of Intelligence Institute (FNii). His current research interests focus on the merging between AI and communication networks. He was selected as the Thomson Reuters Highly Cited researcher and listed in the Worlds' Most Influential Scientific Minds by ScienceWatch in 2014. He was a recipient of the IEEE Signal Processing Society 2012 Best Paper Award. He has served as the general co-chair and the TPC co-chairs for many IEEE conferences. He has been an elected member of IEEE Signal Processing Society SPCOM Technical Committee (2009–2014) and the elected chair of IEEE ComSoc Wireless Technical Committee (2017–2018). He is a member of the Steering Committee of *IEEE Transactions on Big Data* and the chair of the Steering Committee of *IEEE Transactions on Cognitive Communications and Networking*. He has also been serving as an area editor of the *IEEE Signal Processing Magazine* and an associate editor of *IEEE Transactions on Big Data*, *IEEE Transactions on Signal Processing*, *IEEE Journal on Selected Areas in Communications* Series on Green Communications and Networking, and *IEEE Transactions on Wireless Communications*.



Xuemin (Sherman) Shen (Fellow, IEEE) received the PhD degree in electrical engineering from Rutgers University, New Brunswick, NJ, USA, in 1990. He is a University professor with the Department of Electrical and Computer Engineering, University of Waterloo, Canada. His research focuses on network resource management, wireless network security, Internet of Things, AI for networks, and vehicular networks. He is a registered professional engineer of Ontario, Canada, an Engineering Institute of Canada fellow, a Canadian Academy of Engineering fellow,

a Royal Society of Canada fellow, a Chinese Academy of Engineering foreign member, an International fellow of the Engineering Academy of Japan, and a distinguished lecturer of the IEEE Vehicular Technology Society and Communications Society. He received "West Lake Friendship Award" from Zhejiang Province in 2023, President's Excellence in Research from the University of Waterloo in 2022, the Canadian Award for Telecommunications Research from the Canadian Society of Information Theory (CSIT) in 2021, the R.A. Fessenden Award in 2019 from IEEE, Canada, Award of Merit from the Federation of Chinese Canadian Professionals (Ontario) in 2019, James Evans Avant Garde Award in 2018 from the IEEE Vehicular Technology Society, Joseph LoCicero Award in 2015 and Education Award in 2017 from the IEEE Communications Society (ComSoc), and Technical Recognition Award from Wireless Communications Technical Committee (2019) and AHSN Technical Committee (2013). He is the past president of the IEEE Communications Society. He was the vice president for Technical & Educational Activities, vice president for Publications, member-at-large on the Board of Governors, chair of the Distinguished Lecturer Selection Committee, and member of IEEE Fellow Selection Committee of the ComSoc. He served as the editor-in-chief of *IEEE Internet of Things Journal*, *IEEE Network*, and *IET Communications*.



Dong In Kim (Life Fellow, IEEE) received the PhD degree in electrical engineering from the University of Southern California, Los Angeles, CA, USA, in 1990. He was a tenured professor with the School of Engineering Science, Simon Fraser University, Burnaby, BC, Canada. He is currently a distinguished professor with the College of Information and Communication Engineering, Sungkyunkwan University, Suwon, South Korea. He is a fellow of the Korean Academy of Science and Technology and a member of the National Academy of Engineering of Korea.

He was the first recipient of the NRF of Korea Engineering Research Center in Wireless Communications for RF Energy Harvesting from 2014 to 2021. He received several research awards, including the 2023 IEEE ComSoc Best Survey Paper Award and the 2022 IEEE Best Land Transportation Paper Award. He was selected the 2019 recipient of the IEEE ComSoc Joseph LoCicero Award for Exemplary Service to Publications. He was the general chair of the IEEE ICC 2022, Seoul. Since 2001, he has been serving as an editor, an editor at large, and an area editor of the *Wireless Communications I for IEEE Transactions on Communications*. From 2002 to 2011, he served as an editor and a founding area editor of the *Cross-Layer Design and Optimization for IEEE Transactions on Wireless Communications*. From 2008 to 2011, he served as the co-editor-in-chief of the *IEEE/KICS Journal of Communications and Networks*. He served as the founding editor-in-chief of the *IEEE Wireless Communications Letters* from 2012 to 2015. He has been listed as a 2020/2022 Highly Cited researcher by Clarivate Analytics.