

Generative AI as a Service in 6G Edge-Cloud: Generation Task Offloading by In-context Learning

Hao Zhou, Chengming Hu, Dun Yuan, Ye Yuan, Di Wu,

Xue Liu, *Fellow, IEEE*, Zhu Han, *Fellow, IEEE*, and Jianzhong (Charlie) Zhang, *Fellow, IEEE*.

Abstract—Generative artificial intelligence (GAI) is a promising technique towards 6G networks, and generative foundation models such as large language models (LLMs) have attracted considerable interest from academia and industry. This work considers a novel edge-cloud deployment of foundation models in 6G networks. Specifically, it aims to minimize the service delay of foundation models by radio resource allocation and task offloading, i.e., offloading diverse content generation tasks to proper LLMs at the network edge or cloud. In particular, we first introduce the communication system model, i.e., allocating radio resources and calculating link capacity to support generated content transmission, and then we present the LLM inference model to calculate the delay of content generation. After that, we propose a novel in-context learning method to optimize the task offloading decisions. It utilizes LLM’s inference capabilities, and avoids the difficulty of dedicated model training or fine-tuning as in conventional machine learning algorithms. Finally, the simulations demonstrate that the proposed edge-cloud deployment and in-context learning method can achieve satisfactory generation service quality without dedicated model training.

Index Terms—Generative AI, foundation models, 6G edge and cloud, large language models, service delay, in-context learning

I. INTRODUCTION

Generative AI (GAI) has received considerable attention recently, which is capable of analyzing complex data distributions and generating similar new content. Due to its promising features, existing studies have started exploring GAI-enabled 6G networks, e.g., GAI for semantic communication [1] and air-to-ground channel modelling [2]. As a sub-field of GAI, generative foundation models, especially large language models (LLMs), have attracted interest from both academia and industry. For instance, foundation models have been used for network intrusion detection in [3]. Meanwhile, telecom companies have started applying foundation models, e.g., Apple will bring ChatGPT to iPhones with OpenAI, and Qualcomm has developed a mobile platform to support LLMs.

The above progress of academia and telecom industry has demonstrated the great potential of GAI foundation models such as LLMs in 6G networks [1]–[5]. However, despite the advancement, some fundamental and crucial problems are still

not investigated, e.g., practical deployment of GAI foundation models within 6G network architecture, and evaluating the service delay of these generation services in wireless environments. In particular, large GAI models such as GPT-4 usually have billions of parameters, and practical deployment of these large GAI models is critical to support various applications in wireless networks, i.e., network intrusion detection, generation, and management tasks in [3]–[5]. On the other hand, mobile users may send various service requests over wireless channels with diverse preferences, e.g., question-answering tasks need higher accuracy and chatting tasks expect lower delay. Therefore, using one GAI model, i.e., a single LLM, to service the requests of all mobile users is impractical, leading to lower service quality and efficiency.

To this end, we proposed a novel edge-cloud collaboration deployment strategy. We consider small-scale LLMs such as Llama3-8B to be deployed at network edge servers of base stations (BSs), aiming to process tasks efficiently with lower delay, e.g., chatting, information extraction, and content summarization. By contrast, large-scale LLMs, e.g., Llama3-70B and GPT-4, are deployed in the central cloud with abundant computational resources. These large-scale LLMs can generate high-quality content for quality-preferred tasks, such as scientific knowledge, math, and coding-related tasks.

Such an edge-cloud collaboration enables flexible content generation in wireless networks, but it also involves task offloading decisions, i.e., generating content at network edge or offloading tasks to central cloud. Inspired by the recent progress of LLM-based optimization [6], this work further explored in-context learning-based decision-making. Specifically, it uses LLMs to learn from formatted natural language demonstrations and improve the performance on target tasks [7]. Compared with existing machine learning (ML) methods, in-context learning has several advantages: a) Avoiding the complexity of model training and fine-tuning, a well-known bottleneck of conventional ML techniques; b) Following human language instructions to formulate and solve problems, which is far beyond the capabilities of other ML algorithms. In-context learning has been explored to address detection, optimization, prediction tasks, etc [6].

The core contributions of this work are two-fold:

- 1) Firstly, to the best of our knowledge, this work is the first to model the service delay of foundation GAI models in wireless networks, including the communication models for content transmission, and LLM inference models for content generation. It provides a specified metric to evaluate the delay

Hao Zhou, Chengming Hu, Dun Yuan, Ye Yuan, and Xue Liu are with the School of Computer Science, McGill University, Montreal, QC H3A 0E9, Canada. (mails:hao.zhou4, chengming.hu, dun.yuan, ye.yuan3@mail.mcgill.ca, xueliu@cs.mcgill.ca); Di Wu is with the School of Electrical and Computer Engineering, McGill University, Montreal, QC H3A 0E9, Canada. (email: di.wu5@mcgill.ca); Zhu Han is with the Department of Electrical and Computer Engineering at the University of Houston, Houston, TX 77004 USA. (email: hanzhu22@gmail.com); Jianzhong (Charlie) Zhang is with Samsung Research America, Plano, Texas, TX 75023, USA. (email: jianzhong.z@samsung.com).

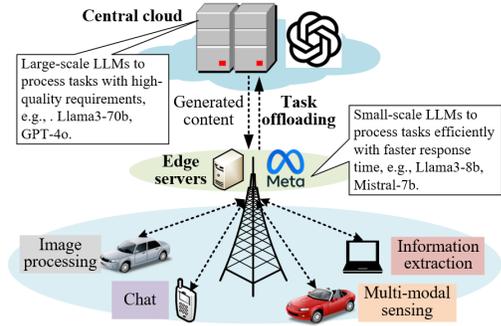


Fig. 1. GAI model deployment and services in wireless networks.

experienced by mobile users. 2) Secondly, we propose a novel in-context learning method for generation task offloading. It avoids the complexity of dedicated model training and fine-tuning as existing ML techniques, using natural language for network management. Finally, the simulations demonstrate that the proposed edge-cloud deployment strategy and in-context learning-based method can achieve satisfactory generation service quality for mobile users.

II. SYSTEM MODEL

Fig. 1 presents the proposed system model, in which mobile users can require various LLM services over wireless networks, e.g., chatting, information extraction and summarization, image generation, etc. Note that these tasks may have different requirements for generated content qualities. Small-scale LLMs are deployed at the network edge to process tasks efficiently with lower delay. By contrast, large-scale LLMs are deployed in the central cloud to handle tasks with high-quality requirements. Therefore, the edge servers must make offloading decisions properly, i.e., processing generation tasks locally or offloading them to the central cloud.

A. Communication Model for Content Transmission

The communication system model considers a downlink transmission for the generated content, involving the wireless transmission delay from BS to end-users, and possible backhaul delay from the network edge to the central cloud as shown in Fig. 1. The transmission delay $t_{k,i}^{trans}$ for downloading the generated content i for user k is

$$t_{k,i}^{tran} = \frac{n_{k,i}^{token} s^{token}}{C_{j,k}} + \alpha_{k,i} t^{back}, \quad (1)$$

where $n_{k,i}^{token}$ is the LLM output token numbers for input prompt i , $C_{j,k}$ is the link capacity for download transmission between user k and BS j , s^{token} is the byte size per token [8]. $\frac{n_{k,i}^{token} s^{token}}{C_{j,k}}$ represents the transmission delay from BSs to users. t^{back} is the backhaul delay caused by task offloading. $\alpha_{k,i}$ is the task offloading decision: $\alpha_{k,i} = 1$ means offloading tasks to central cloud, while $\alpha_{k,i} = 0$ indicates network edge implementation. We assume a fixed t^{back} in (1), which is a setting used in many task offloading-related studies [9]. The link capacities $C_{j,k}$ can be calculated by

$$C_{j,k} = \sum_{q \in \mathcal{Q}_j} b_q \log\left(1 + \frac{p_{j,q} g_{j,q,k} z_{j,q,k}}{\sum_{j' \in \mathcal{J}_{-j}} p_{j',q'} g_{j',q',k'} z_{j',q',k'} + b_q N_0}\right), \quad (2)$$

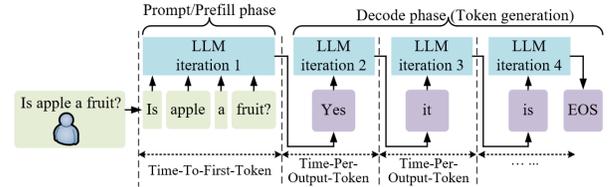


Fig. 2. LLM inference process illustration. (EOS: end-of-sequence).

where \mathcal{Q}_j is the resource blocks (RBs) set of the j^{th} BS, b_q is the bandwidth of the q^{th} RB, $p_{j,q}$ is the transmission power of the q^{th} RB, $g_{j,q,k}$ is the channel gain between BS and user, $z_{j,q,k}$ is a binary indicator to represent whether the q^{th} RB is allocated to user k , and N_0 is the noise power density. \mathcal{J}_{-j} is the set of BSs except j^{th} BS, and $\sum_{j' \in \mathcal{J}_{-j}} p_{j',q'} g_{j',q',k'} z_{j',q',k'}$ is inter-cell interference. We assume orthogonal frequency-division multiplexing is deployed to avoid intra-cell interference.

B. LLM Inference Model for Content Generation

As shown in Fig. 2, the LLM inference process mainly consists of prefill phase and decode phase [10]. The user question is split into smaller tokens, then the LLM processes the input tokens as a next-token predictor by autoregressive decoding. Fig. 2 demonstrates that the LLM inference time can be generally divided into two parts: Time-To-First-Token (TTFT) refers to the time to generate the first token, and Time-Per-Output-Token (TPOT) indicates the time to generate each following token [10]. Therefore, the total generation time for a prompt i from user k is

$$t_{k,i}^{gen} = t^{TTFT} + n_{k,i}^{token} t^{TPOT}, \quad (3)$$

where $t_{k,i}^{gen}$ is the total generation time, t^{TTFT} is TTFT time, $n_{k,i}^{token}$ is the number of tokens generated for prompt i , and t^{TPOT} is the TPOT time. LLM is a complicated system with a huge number of parameters, and t^{TTFT} and t^{TPOT} are affected by many factors, e.g., model architecture, hardware constraints, and task types. Therefore, it is extremely difficult to calculate the exact generation time for each task. However, (3) provides a practical approach to quantify LLM generation time since the TTFT and TPOT values of many LLMs can be easily tested and obtained for evaluation purposes [11].

C. Problem Formulation

This work aims to minimize the total generation and transmission delay of all K users from BS j in wireless networks, and meanwhile satisfy the quality requirements of the generated content. The problem formulation is defined as

$$\min_{\substack{\alpha_{k,i}, \\ z_{j,q,k}}} \sum_{k=1}^K \sum_{i=1}^{I_k} \left(t_{k,i}^{tran} + \alpha_{k,i} t_{k,i}^{cloud} + (1 - \alpha_{k,i}) t_{k,i}^{edge} \right) \quad (4)$$

$$\text{s.t.} \quad (1) \text{ and } (2), \quad (4a)$$

$$z_{j,q,k}, \alpha_{k,i} \in \{0, 1\}, \quad (4b)$$

$$t_{k,i}^{edge} = t^{TTFT,edge} + n_{k,i}^{token} t^{TPOT,edge}, \quad (4c)$$

$$t_{k,i}^{cloud} = t^{TTFT,cloud} + n_{k,i}^{token} t^{TPOT,cloud}, \quad (4d)$$

$$\tau_{k,i} \leq \alpha_{k,i} \tau^{cloud} + (1 - \alpha_{k,i}) \tau^{edge}, \quad (4e)$$

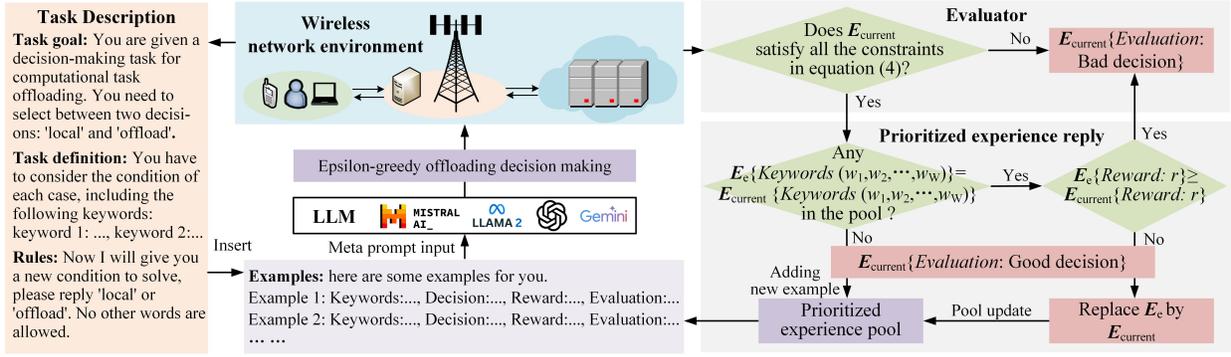


Fig. 3. LLM-enabled in-context learning for task offloading.

where $\alpha_{k,i}$ is the task offloading decision, $z_{j,q,k}$ is the RB allocation decision defined in (2), and I_k is the total number of service requests from user k . We assume a small-scale LLM is deployed at the network edge as in constraint (4c), in which $t_{k,i}^{edge}$, $t_{TTFT,edge}$, and $t_{TPOT,edge}$ represent total generation time, TTFT, and TPOT values. Similarly, $t_{k,i}^{cloud}$, $t_{TTFT,cloud}$, and $t_{TPOT,cloud}$ are defined for the cloud LLM in (4d). In addition, the quality of the generated content is also a crucial metric to evaluate generation services. (4e) is the generation quality constraint, indicating that the quality index of selected edge LLM τ^{edge} or cloud LLM τ^{cloud} should be higher than the user requirement $\tau_{k,i}^1$.

Finally, (4) includes both radio resource allocation and task offloading decisions, in which $z_{j,q,k}$ for radio resource allocation and $\alpha_{k,i}$ for task offloading. For radio resource allocation, we apply a classic proportional fairness algorithm because: 1) proportional fairness is a practical method that has been widely used in many existing studies; 2) this work aims to understand generative foundation model service properties, and it is reasonable to apply a well-known resource allocation method to better focus on foundation models. For task offloading, we propose an in-context learning-based task offloading approach in the following Section III.

III. IN-CONTEXT LEARNING FOR TASK OFFLOADING

In-context learning refers to the process of learning from formatted natural language-based task descriptions and examples, aiming to improve the performance of target tasks [7]. Considering a query input x and possible candidate answers $\mathcal{Y} = \{y_1, y_2, \dots, y_{|\mathcal{Y}|}\}$, a set of examples are provided as $\mathcal{E} = \{E_1, E_2, \dots, E_{|\mathcal{E}|}\}$, in which each $E_e \in \mathcal{E}$ consists of input-output pairs as $E_e = (x_e, y_e)$. The probability of generating a specific output y^* is

$$Pr(y^*|x) \triangleq f_{LLM}(x, y^*, \{E_1, E_2, \dots, E_{|\mathcal{E}|}\}, D), \quad (5)$$

where $f_{LLM}(\cdot)$ is a scoring function and D is the task description. Then the final output answer \hat{y} is the candidate answer with the highest probability

$$\hat{y} = \arg \max_{y \in \mathcal{Y}} (Pr(y|x)). \quad (6)$$

¹The generation quality index can be obtained by testing LLMs on task datasets, e.g., Chatbot Arena, Multi-task Language Understanding, etc [11].

The above (5) and (6) prove that the output \hat{y} depends on the input x , task description D , and example set \mathcal{E} . In this work, LLM outputs \hat{y} refers to the decision between local implementation and offloading, the input x involves service types and the estimated output token size, D is the offloading task description, and \mathcal{E} is a set of previous examples, which will be introduced in following subsections.

A. Prompting System Design

The overall organization of the proposed LLM-enabled in-context learning is shown in Fig. 3 with the following steps:

Step 1: Task description. Based on network environments, the task description D in (5) is first defined by

$$D = \{\text{Task goal}, \text{Task definition}, \text{Rules}\}, \quad (7)$$

in which the “Task goal” specifies the target problem with two decision variables “local” and “offload”. The “Task definition” indicates the status variables affecting offloading decisions. Here we consider “Service types” and “Estimated output token size”, which also means the input x in (5) is

$$x = \{\text{Service types}, \text{Estimated output token size}\}. \quad (8)$$

Additionally, extra “Rules” are applied to LLMs, e.g., replying “local” or “offload” only to improve the output accuracy.

Step 2: Example design. The task description D will be combined with the example set $\mathcal{E} = \{E_1, E_2, \dots, E_{|\mathcal{E}|}\}$ as a meta prompt input to the LLM, producing task offloading decisions α . The example $E_e \in \mathcal{E}$ is defined by

$$E_e \{\text{Keywords} : (w_1, w_2, \dots, w_W), \text{Decision} : \text{Local/Offload}, \text{Reward} : r, \text{Evaluation} : \text{Good/Bad decision}\} \quad (9)$$

in which the “Keywords” refer to the values of “Service types” and “Estimated output token size” defined in (8). Inspired by reinforcement learning, a reward metric is defined to evaluate the system performance by jointly considering service delay and quality requirements.

$$r = T^{\text{Target}} - t^{\text{total}} - r^{\text{penalty}} \quad (10)$$

where T^{Target} is the target delay, and t^{total} is the objective function defined in (4) and $t^{\text{total}} = t_{k,i}^{\text{tran}} - \alpha_{k,i} t_{k,i}^{\text{cloud}} - (1 - \alpha_{k,i}) t_{k,i}^{\text{edge}}$. Here r^{penalty} is a penalty item, aiming to balance the service delay and generated content quality. r^{penalty} is

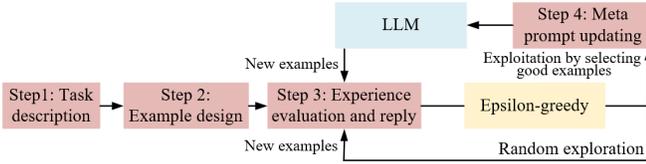


Fig. 4. The overall optimization procedure of the proposed technique

a preset positive value if the constraint in (4) is violated. A high $r^{penalty}$ value means strict requirements on generated content quality, while a lower $r^{penalty}$ value allows lower generation quality to achieve lower service delay. Otherwise, if all constraints are satisfied in (4), $r^{penalty} = 0$. Such a definition is also a widely used approach in reinforcement learning studies to handle optimization constraints. Therefore, (10) provides a comprehensive metric to minimize the total delay under the constraints.

Step 3: Experience evaluation and replay. Given the offloading decision, the current network operation results become a new experience example $E_{current}$, which will be sent to the evaluator module as in Fig. 3. Specifically, $E_{current}$ is considered a bad decision if the constraints in (4) are not satisfied; otherwise, $E_{current}$ is sent to the prioritized experience replay module for further evaluation, which will be introduced in the following Subsection III-B.

Step 4: Meta prompt updating. The experience pool in Step 3 will serve as a new example set \mathcal{E}_{next} , and the LLM will generate the next decision using the same task description D and the updated example set \mathcal{E}_{next} . Steps 3 and 4 will be repeated until algorithm performance converges, which will be analyzed in the following Fig. 4.

B. Prioritized Experience Replay and Exploration Strategies

The above Steps 3 and 4 show that properly updating the experience pool is crucial, and this subsection will present two key techniques, namely prioritized experience replay and epsilon-greedy exploration, aiming to identify the most useful examples for experience replay. In particular, Fig. 3 shows that the prioritized experience replay includes two rules:

1) If $E_{current}$ is a new example, which means that $E_{current}\{\text{Keywords}:(w_1, w_2, \dots, w_{\mathcal{W}})\} \neq E_e\{\text{Keywords}:(w_1, w_2, \dots, w_{\mathcal{W}})\}$ for $\forall E_e \in \mathcal{E}$, then $E_{current}$ is always considered as a “Good decision” since there is no existing example of this condition in the current experience pool \mathcal{E} .

2) If $\exists E_e \in \mathcal{E}$ that has the same keyword values as $E_{current}$ with $E_e\{\text{Keywords}:(w_1, w_2, \dots, w_{\mathcal{W}})\} = E_{current}\{\text{Keywords}:(w_1, w_2, \dots, w_{\mathcal{W}})\}$, then we will compare their reward values. If E_e has a higher reward than $E_{current}$, then $E_{current}$ is considered a “Bad decision” and the experience pool \mathcal{E} remains unchanged. Otherwise, if $E_{current}$ has a higher reward, then $E_{current}$ becomes a better example as a “Good decision”. E_e will be replaced by $E_{current}$ in the experience pool, and a new example set \mathcal{E}_{next} is generated.

Fig. 4 summarizes the overall procedure of the proposed techniques. We apply the well-known epsilon-greedy policy to balance exploration and exploitation, which is a fundamental problem in many optimization tasks. Specifically, actions are

randomly selected with probability ϵ ; otherwise, LLM will make decisions. Combining the epsilon-greedy policy with the proposed experience replay technique can send the best examples found to LLMs, and also explore new examples to improve the experience pool. Therefore, with plenty of explorations, the examples in the experience pool will be constantly improved, and the LLM output will converge if no better examples can be found².

Finally, we present 3 baseline algorithms. Baseline 1: Latest experience-based in-context learning, using the latest experience as examples in the prompt. Baseline 2: In-context learning without exploration, and all decisions are made by LLMs. Baseline 3: We consider deep reinforcement learning (DRL) as an optimal baseline since DRL techniques have been very widely applied to solve various network optimization problems, in which the state is defined by (8), the action is offloading decision, and the reward is shown as (10).

IV. PERFORMANCE EVALUATION

A. Simulation Settings

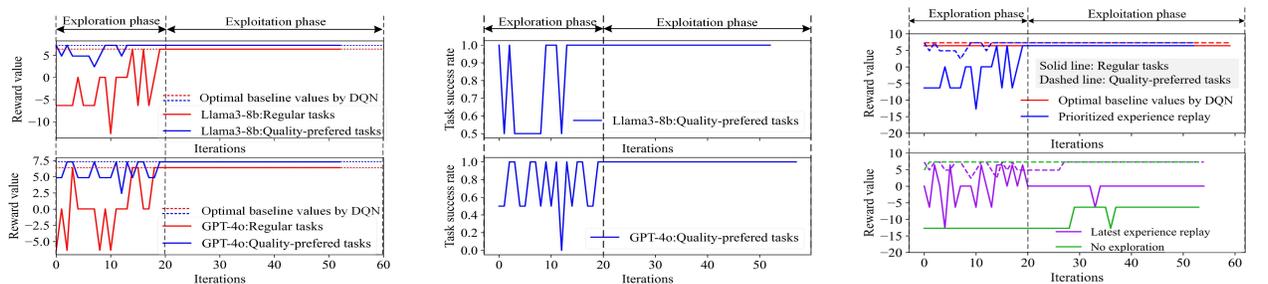
We assume 20 users are randomly distributed in a cell based on 3GPP urban networks. The mobile users have two kinds of generation tasks: regular tasks such as chatting, translation, and summarization, and quality-preferred tasks such as scientific knowledge reasoning and coding. Such quality-based task classification aligns well with the evaluation schemes for LLMs, e.g., Chatbot Arena and Multi-task Language Understanding (MMLU). The average output size is 1,000 tokens, and each token corresponds to about 4 bytes according to an OpenAI report [8]. We evaluate 2 LLMs for decision-making: Llama3-8B as a small-scale model that can be deployed at the network edge and GPT-4o as the latest large-scale LLM model as a benchmark for LLM-enabled methods. For the network settings, we assume the edge LLM’s TTFT and TPOT values are 0.23 and 1/75 second, while the values for the cloud LLM are 0.42 and 1/32 [11].

B. Simulation Results

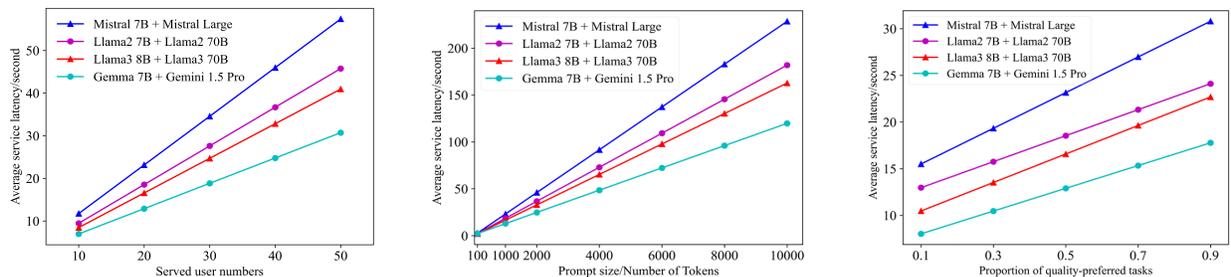
Fig. 5(a) shows the system reward of different tasks using Llama3-8B and GPT-4o, in which all tasks converge to a stable reward after exploration. Here, the reward values indicate the overall service performance as defined in (10). Llama3-8B and GPT-4o also achieve comparable performance as conventional DRL-based optimal baseline, demonstrating the potential of LLM-based optimization techniques. In addition, Fig. 5(b) presents the service success rate for quality-preferred tasks, which means that the service score constraint in (4) should be fulfilled. It shows that the proposed in-context learning method can offload tasks properly to satisfy user requirements.

We further compare the experience replay and exploration methods. in Fig. 5(c). It shows that the proposed prioritized experience replay method can obtain comparable performance as optimal DRL baseline values. Specifically, epsilon-greedy

²Similar assumptions can be found in the convergence analyses of classic tabular-based Q-learning, e.g., “Q-learning is guaranteed to converge when visiting each state-action tuple infinitely times”.



(a) System reward comparison of different tasks using Llama3-8B and GPT-4o. (b) Service success rate of quality-preferred tasks using using Llama3-8B and GPT-4o. (c) Comparisons under different experience replay and exploration methods using Llama3-8B.



(d) Average service delay with different user numbers using different LLM combinations. (e) Average service delay with different prompt size using different LLM combinations. (f) Average service delay with different tasks proportions using different LLM combinations.

Fig. 5. Simulation results and comparisons

exploration can try different decisions and collect good examples. These good examples will then be used in prioritized experience replay to guide the future decisions of LLMs. By contrast, latest experience reply and no-exploration methods present lower rewards as shown in the lower part of Fig. 5(c).

Finally, Figs. 5(d), 5(e) and 5(f) present the service delay of various LLMs under different user numbers, size of output tokens, and task distributions. In particular, when the proportion of quality-preferred tasks increases, these generation tasks must be offloaded to large-scale LLMs at cloud, leading to higher service delays. Meanwhile, Gemma 7B + Gemini 1.5 Pro can achieve the lowest overall delay than other combinations. This is because of the low TTFT values of these two LLMs, which are 1/155 second for Gemma 7B and 1/58 second for Gemini 1.5 Pro. Other LLMs have much higher values, i.e., 1/89 second for Llama2-7B and 1/40 second for Llama2-70B [11]. In addition, some supplementary experiment results and analyses can be found in [12].

V. CONCLUSION

GAI is a promising technique for future 6G networks, and this work investigates the generation task offloading problems in 6G edge-cloud. It proposes a novel in-context learning method for generation task offloading, and the simulations demonstrated that the proposed technique can achieve satisfactory generation service quality. It is worth noting that such an edge-cloud deployment method may also increase the vulnerabilities of LLMs, e.g., inference and extraction attacks. In the future, we will explore secure and robust deployment strategies for LLMs in 6G networks, better guaranteeing the generation service reliability and quality.

REFERENCES

- [1] L. Qiao, M. B. Mashhadi, Z. Gao, C. H. Foh, P. Xiao, and M. Bennis, "Latency-Aware Generative Semantic Communications with Pre-Trained Diffusion Models," *IEEE Wireless Communications Letters*, 2024.
- [2] A. Giuliani, R. Nikbakht, G. Geraci, S. Kang, A. Lozano, and S. Rangan, "Spatially Consistent Air-to-Ground Channel Modeling via Generative Neural Networks," *IEEE Wireless Communications Letters*, 2024.
- [3] M. Fu, P. Wang, M. Liu, Z. Zhang, and X. Zhou, "IoV-BERT-IDS: Hybrid Network Intrusion Detection System in IoV Using Large Language Models," *IEEE Transactions on Vehicular Technology*, 2024.
- [4] Q. Liu, J. Mu, D. ChenZhang, Y. Liu, and T. Hong, "LLM Enhanced Reconfigurable Intelligent Surface for Energy-Efficient and Reliable 6G IoV," *IEEE Transactions on Vehicular Technology*, 2024.
- [5] S. Javaid, R. A. Khalil, N. Saeed, and et al., "Leveraging Large Language Models for Integrated Satellite-Aerial-Terrestrial Networks: Recent Advances and Future Directions," *arXiv:2407.04581*, 2024.
- [6] H. Zhou, C. Hu, Y. Yuan, Y. Cui, Y. Jin, and et al., "Large language model (LLM) for telecommunications: A comprehensive survey on principles, key techniques, and opportunities," *arXiv preprint arXiv:2405.10825*, 2024.
- [7] Q. Dong, L. Li, D. Dai, C. Zheng, Z. Wu, and et al., "A survey on in-context learning," *arXiv preprint arXiv:2301.00234*, 2022.
- [8] OpenAI, "What are tokens and how to count them?" 2024. [Online]. Available: <https://help.openai.com/en/articles/4936856-what-are-tokens-and-how-to-count-them>
- [9] E. El Haber, T. M. Nguyen, C. Assi, and W. Ajib, "Macro-cell assisted task offloading in MEC-based heterogeneous networks with wireless backhaul," *IEEE Transactions on Network and Service Management*, vol. 16, no. 4, pp. 1754–1767, Sep. 2019.
- [10] P. Patel, E. Choukse, C. Zhang, Í. Goiri, A. Shah, S. Maleki, and R. Bianchini, "Splitwise: Efficient generative LLM inference using phase splitting," *arXiv preprint arXiv:2311.18677*, 2023.
- [11] "Artificial analysis," 2024. [Online]. Available: <https://artificialanalysis.ai>
- [12] "Supplementary materials of "Generative AI as a Service in 6G Edge-Cloud: Generation Task Offloading by In-context Learning";", 2024. [Online]. Available: <https://github.com/haozhou1995/Supplementary-Experiments-and-Discussions.git>