

Tool-Aided Evolutionary LLM for Generative Policy Toward Efficient Resource Management in Wireless Federated Learning

Chongyang Tan¹, Ruoqi Wen¹, Rongpeng Li¹, *Senior Member, IEEE*, Zhifeng Zhao², *Member, IEEE*, Ekram Hossain³, *Fellow, IEEE*, and Honggang Zhang⁴, *Fellow, IEEE*

Abstract—Federated Learning (FL) enables distributed model training across edge devices in a privacy-friendly manner. However, its efficiency heavily depends on effective device selection and high-dimensional resource allocation in dynamic and heterogeneous wireless environments. Conventional methods demand a confluence of domain-specific expertise, extensive hyperparameter tuning, and/or heavy interaction cost. This paper proposes a Tool-aided Evolutionary Large Language Model (T-ELLM) framework to generate a qualified policy for device selection in a wireless FL environment. Unlike conventional optimization methods, T-ELLM leverages natural language-based scenario prompts to enhance generalization across varying network conditions. The framework decouples the joint optimization problem mathematically, enabling tractable learning of device selection policies while delegating resource allocation to convex optimization tools. To facilitate the evolutionary process, T-ELLM interacts with a sample-efficient, model-based virtual learning environment that captures the relationship between device selection and learning performance. This developed virtual environment reduces reliance on real-world interactions, thus minimizing communication overhead while refining the LLM-based decision-making policy through group relative policy optimization. Theoretical analysis proves that the discrepancy between virtual and real environments is bounded, ensuring the advantage function learned in the virtual environment maintains a provably small deviation from real-world conditions. Experimental results demonstrate that T-ELLM outperforms benchmark methods in energy efficiency and exhibits robust adaptability to environmental changes.

Index Terms—Large language model, generative policy, wireless federated learning, resource management, convex optimization, reinforcement learning.

Received 15 May 2025; revised 14 October 2025; accepted 9 December 2025. Date of publication 18 December 2025; date of current version 24 February 2026. This work was supported in part by the National Key Research and Development Program of China under Grant 2024YFE0200600, in part by Zhejiang Provincial Natural Science Foundation of China under Grant LR23F010005, and in part by Huawei Cooperation Project under Grant TC20240829036. (*Corresponding author: Rongpeng Li.*)

Chongyang Tan, Ruoqi Wen, and Rongpeng Li are with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310058, China (e-mail: cyatan@zju.edu.cn; wenruoqi@zju.edu.cn; lirongpeng@zju.edu.cn).

Zhifeng Zhao is with Zhejiang Laboratory, Hangzhou 311121, China, and also with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310058, China (e-mail: zhaozf@zhejianglab.org).

Ekram Hossain is with the University of Manitoba, Winnipeg, MB R3T 5V6, Canada (e-mail: ekram.hossain@umanitoba.ca).

Honggang Zhang is with Macau University of Science and Technology, Macau 999078, China (e-mail: hg Zhang@must.edu.mo).

Digital Object Identifier 10.1109/JSAC.2025.3645754

I. INTRODUCTION

NEXT-GENERATION (xG) wireless communication systems are envisioned to support various intelligent applications and services [1], empowered by the exponential growth of wireless edge devices, such as mobile phones and sensors. As a prominent paradigm [2], [3], Federated Learning (FL) emerges by facilitating collaborative model training across decentralized devices while maintaining data locality in a privacy-friendly manner. Nevertheless, the deployment of FL in wireless networks faces severe challenges, arising from the underlying heterogeneous computation and communication capabilities across devices [4], [5]. Correspondingly, a proper resource management policy for FL, which selects the participating devices and calibrates the utilized communications and computing resources, will need to be developed in dynamic wireless environments [6]. Typically, such a problem results in a high-dimensional optimization problem that can be *partially* solved by heuristic solutions [7], [8] or Reinforcement Learning (RL)-based approaches [9], [10], [11], [12], [13]. However, heuristic solutions often demand a confluence of domain-specific expertise and extensive tuning before adapting to unseen scenarios, while RL-based approaches require heavy interactions with the environment and become sluggish for large system dimensions and changing system dynamics [14], [15], [16]. Therefore, there is a strong incentive to find alternative solutions [17].

A. Related Works

To implement FL over wireless networks, in each training round, edge devices upload the locally trained updates to a centralized server via wireless links, in exchange for aggregated models. Through several training rounds, the performance of the eventually learned global model is primarily impacted by resource and data heterogeneity. For example, disparities in computational and communication capabilities across devices lead to imbalanced time and energy consumption, while varying dataset sizes and non-Independent and Identically Distributed (non-IID) data distributions can result in biased model updates. Both would degrade the learning effectiveness [4], [5]. Additionally, the deployment of FL encounters constraints due to the energy and time [8] budget [10], [18]. For instance, the authors in [19] and [20]

formulate the joint learning and communication problem as the minimization of total energy consumption under latency constraints and provide computationally efficient closed-form solutions to optimize critical resources, including bandwidth allocation, CPU frequency, and transmission power. But, full device participation is bluntly assumed in [19] and [20]. However, mobilizing the participation of a subset of qualified edge devices in each training round can contribute to improving learning efficiency. For example, biasing client selection towards clients with higher local loss [21] proves to yield faster convergence than the random selection in Federated Averaging (FedAvg) [2]. Also, a timer can be heuristically set to avoid the participation of straggler devices [8]. In [7], a guided participant selection scheme improves FL performance by jointly optimizing system and data utility through an epsilon-greedy algorithm. Moreover, a joint device scheduling and bandwidth allocation policy, which maximizes model accuracy within a fixed training time budget by a hybrid greedy-optimization approach, is proposed in [22], while energy constraints are neglected therein. More importantly, adapting these heuristics to unexpected scenarios, such as tuning hyperparameters to approximate FL performance limits [8], [22] or setting appropriate timers [7] when system bandwidth or device resources (e.g., maximum transmission power, dataset size) change, typically demands iterative adjustments and domain-specific expertise.

RL provides an alternative methodology to solve wireless FL problems. Through continuously interacting with the environment, RL agents evaluate candidate policies in the wireless FL environment and gradually learn appropriate real-time policies. For example, [12] develops a Deep RL (DRL) method to select devices to upload their local model for aggregation, to maximize validation accuracy. Nevertheless, it overlooks energy and time constraints and assumes all devices participate in local training. Similarly, [11] adopts a Q-learning-based method to identify near-optimal participating devices. Furthermore, to ameliorate the impact of high dimensionality on single-agent RL with large-scale devices, [9] introduces a Multi-Agent Reinforcement Learning (MARL) framework that jointly optimizes model accuracy, processing latency, and communication efficiency. On the other hand, [13] improves the energy efficiency of FL by managing the CPU-cycle frequency of mobile devices based on DRL. To further improve the efficiency of FL, the works in [23], [24], and [25] leverage DRL-based approaches to jointly optimize the device selection and bandwidth allocation. However, the above DRL-based methods only optimize partial parameters and rely on extensive training in a given environment. Therefore, the changes in environment or data often imply cumbersome retraining or even architectural redesign. Unfortunately, a real-world wireless scenario constantly evolves, and it inevitably induces substantial engineering cost. Meanwhile, direct interaction with the environment incurs substantial communication overhead. Although imitation learning [15] can acquire a policy by learning from the optimal trajectories to reduce interaction overhead, there are no such expert systems strictly in this FL scenario. Besides, the training difficulty, which increases exponentially along with the potentially involved

devices and managed resources, limits the practicability [14], [16].

The remarkable capabilities of transformer-based Large Language Models (LLMs) [26], [27], [28], [29] demonstrate the feasibility of a single general-purpose model, trained on extensive text corpora, for diverse, generalizable task accomplishment [30]. Capitalizing on the success of LLMs, we resort to an LLM-driven textual description approach to leverage its inherent flexibility and generalization. Leveraging fine-tuned LLMs for FL optimization starkly contrasts with those efforts, which focus on enhancing the training capabilities of LLMs in wireless FL [31], [32]. However, directly applying standard LLM methodologies to this specialized domain presents several fundamental obstacles. Prominently, common approaches like training and fine-tuning require vast real-world interaction datasets, which are prohibitively time-consuming and costly to collect [33], [34]. The non-availability of expert decision-making data for FL efficiency can add to the difficulty of reliably teaching an LLM to interpret wireless FL states and generate effective actions. Furthermore, general-purpose LLMs are inherently weak at precise mathematical reasoning and can be prone to hallucination [35], making them ill-suited for optimization-centric tasks, such as resource management in wireless FL. Meanwhile, the high latency associated with processing overlong token sequences for complex reasoning makes the sole adoption of LLM infeasible in time-sensitive wireless FL systems [36]. Therefore, due to the domain-specific expertise and computationally intensive optimization operations, extra effort is still needed to make LLMs qualified for decision-making in wireless FL.

B. Contributions

In this paper, we propose a Tool-aided Evolutionary LLM (T-ELLM) approach to realize efficient device selection and resource allocation in wireless FL. Specifically, T-ELLM adopts a natural language-based scenario prompt, where linguistic flexibility helps to improve the generalization capability [37]. Afterward, on top of a mathematics-established problem decoupling, T-ELLM capably yields device selection and resource management results by answering the scenario query and invoking convex optimization tools, respectively. Furthermore, to better tackle possible environmental changes, alongside available convex optimization tools, T-ELLM builds a sample-efficient model-based virtual learning environment, which characterizes the relationship between learning effectiveness and device selection. Once this high-fidelity model is established, the T-ELLM agent can be trained for a vast number of episodes entirely offline, at virtually zero communication cost. In other words, since the interaction cost is entirely induced by collecting samples for the simpler virtual environment modeling, this model-based approach yields high sample efficiency by drastically reducing the communication and time overhead for practically interacting LLMs with the real world. The main contributions of this paper are summarized as follows:

- We consider a high-dimensional decision-making problem in wireless FL scenarios, and adopt a tool-based

evolutionary LLM approach, which is referred to as T-ELLM. In particular, we fine-tune an LLM, which can respond to the linguistic description of wireless FL scenarios with an appropriate output of device selection. The astonishing contextual understanding capability in LLM empowers T-ELLM to cope with environmental changes flexibly, significantly reducing manual design effort and improving adaptability.

- We establish the mathematical rationale to decouple the joint device selection and resource management problem, which significantly reduces the decision-making space to be learned. On this basis, augmented with contextual convex optimization-based resource management tools, T-ELLM invokes and adapts tools to perform joint optimization.
- We incorporate a sample-efficient model-based virtual learning environment to ground the LLM in wireless FL scenarios. In combination with the resource management tool, the model-based approach enables the LLM to learn from a simulated environment by Group Relative Policy Optimization (GRPO) [38], due to its merits in efficient trajectory collection. The continual learning in a high-fidelity simulation environment improves the decision-making capabilities of LLM, while significantly reducing communication overhead from practical interactions in the real world.
- We prove that the overall discrepancy between the virtual and real environment is provably bounded, yielding a limited deviation for the advantage function learned from the virtual environment from its counterpart in the real environment. Besides, the experimental results show that the proposed framework outperforms other benchmarks in terms of energy consumption, while demonstrating its adaptability to environmental changes.

C. Paper Structure

The remainder of this paper is organized as follows. We introduce the system model and problem formulation in Section II. Afterward, Section III presents how to decouple the optimization problem and explains the rationale behind it. In Section IV and Section V, we propose the T-ELLM scheme and analyze its convergence properties. Afterward, we present the simulation results in Section V. Finally, we conclude the paper in Section VII.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we first present the FL model, followed by the definition of a wireless communication system model. Finally, we formulate the problem of efficient wireless FL resource management.

A. System Model

1) *Wireless FL Model*: We consider a wireless FL system consisting of a Base Station (BS) equipped with a central server and N distributed devices represented by $\mathcal{N} = \{0, 1, \dots, n, \dots, N-1\}$. Each device n has its local dataset, denoted by D_n , with a size $|D_n|$. For a given Neural Network

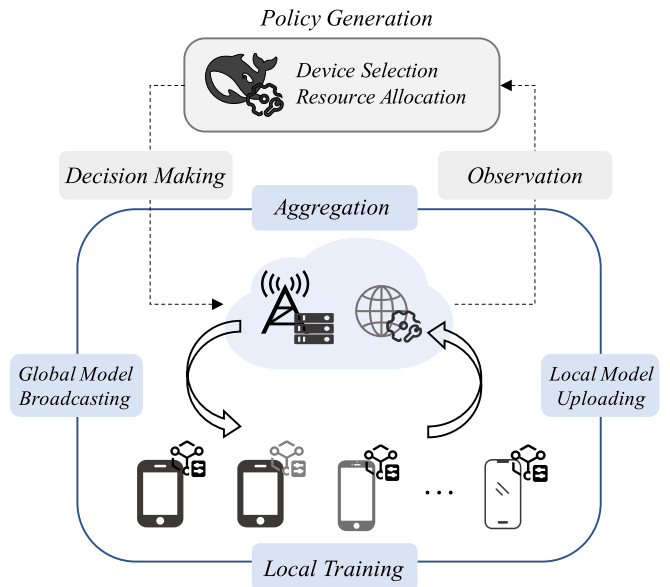


Fig. 1. An illustration of the wireless FL environment with a device selection and resource management policy generator.

(NN) training task, define the loss function as $l(\omega, x)$, where ω represents the parameters of the NN and x is a sample in the dataset. Subsequently, the local training loss of device n can be calculated as:

$$F_n(\omega_n) = \frac{1}{|D_n|} \sum_{x \in D_n} l(\omega_n, x), \quad (1)$$

where ω_n is the local parameters of device n . While FL aims to find global parameters ω_G such that the loss is minimized across the entire dataset $D = \bigcup_{n \in \mathcal{N}} D_n$ without sharing any datasets, the corresponding optimization problem can be formulated as:

$$\omega_G^* = \arg \min_{\omega_G} F(\omega_G), \quad (2)$$

where $F(\omega_G) = \sum_{n \in \mathcal{N}} \frac{|D_n|}{|D|} F_n(\omega_G)$ is the global loss function.

As illustrated in Fig. 1, FL algorithms typically solve (2) through an iterative three-step process, which encompasses local training, aggregation, and synchronization, repeated over multiple rounds. Taking the example of *FedAvg* [2], in the t -th communication round, the central server broadcasts the $(t-1)$ -th global model parameters ω_G^{t-1} to each device, where ω_G^0 indicates an initialized model. After getting the global model, each device n uses its local dataset to update the parameters of the local model I times, i.e.,

$$\omega_n^t(i) = \omega_n^t(i-1) - \eta \nabla F_n(\omega_n^t(i-1)), \quad (3)$$

where $i \in [1, I]$ represents the index of local updates and $\omega_n^t(0) = \omega_G^{t-1}$. It is worth noting that only a portion of the devices truly participate in the training, and these selected devices in the t -th round are denoted as $\mathbb{S}_t \subset \mathcal{N}$. All these selected devices transmit the trained local models to the central server, in exchange for a new, aggregated global model, i.e.,

$$\omega_G^t = \frac{1}{|D|} \sum_{n \in \mathbb{S}_t} |D_n| \omega_n^t, \quad (4)$$

where $\omega_n^t = \omega_n^t(I)$. Considering the impact of the subset \mathbb{S}_t on the convergence behavior of FL [5], the central server shall calibrate a means to determine device participation, thereby optimizing training performance.

2) *Computation and Communication Model for FL*: We consider the energy and time consumption associated with the computation and communication processes in the wireless FL system. Due to the high transmission power of the BS and sufficient downlink bandwidth for global model broadcast, the consumed time and energy at the BS are assumed to be negligible [39], [40]. Therefore, we only focus on the device computation and communication.

- *Local Computation* [23]: Let $f_{t,n}$ denote the CPU frequency (in CPU cycles per second) of device n in the t -th round. The time required for I times of local model training on device n in the t -th round can be expressed as $T_{t,n}^{\text{comp}} = \frac{C|D_n|I}{f_{t,n}}$, where C (cycles per sample) represents the number of CPU cycles needed to process one sample by the backpropagation algorithm. Additionally, the energy consumption of device n for local model training in the t -th round is given by $E_{t,n}^{\text{comp}} = \zeta C|D_n|I f_{t,n}^2$, where ζ is the effective switched capacitance that depends on the chip architecture.
- *Model Transmission*: After completing local model training, the selected devices upload their local model parameters to the BS via Frequency Division Multiple Access (FDMA). The achievable transmission rate of local device n in the t -th round is given by $v_{t,n} = B_{t,n} \log_2 \left(1 + \frac{p_{t,n} G_{t,n}}{N_0 B_{t,n}} \right)$, where $B_{t,n}$ and $p_{t,n}$ denotes the bandwidth and transmission power allocated to device n in the t -th communication round respectively, $G_{t,n}$ is the channel coefficient of device n in the t -th round under the current environmental conditions, and N_0 is the noise power spectral density. Let s_n denote the size (in bits) of the model parameters transmitted by device n . Therefore, the time required for device n to upload its model parameters in the t -th communication round is given by $T_{t,n}^{\text{com}} = \frac{s_n}{v_{t,n}}$. Subsequently, the energy consumption of device n in the t -th round for uploading model parameters can be expressed as $E_{t,n}^{\text{com}} = p_{t,n} T_{t,n}^{\text{com}} = \frac{p_{t,n} s_n}{v_{t,n}}$.

Above all, for each communication round, the total energy consumption is modeled as:

$$E_t = \sum_{n \in \mathbb{S}_t} (E_{t,n}^{\text{com}} + E_{t,n}^{\text{comp}}). \quad (5)$$

Similarly, the total time consumption is given by

$$T_t = \max_{n \in \mathbb{S}_t} \{T_{t,n}^{\text{com}} + T_{t,n}^{\text{comp}}\}. \quad (6)$$

B. Problem Formulation

We tackle the joint optimization of device selection and resource allocation in wireless FL. In particular, we aim to maximize energy efficiency while adhering to diverse computational, communication, and task-specific constraints. For example, some scenarios prioritize minimizing runtime, while others only require meeting time-related Quality-of-Service (QoS). Similarly, bandwidth allocation may be equal or flexible, contingent on the scenario. First, we formulate an

objective function that balances FL learning performance and device energy consumption, expressed as:

$$r(\Xi_t, E_t) = (1 - \sigma) \frac{\Xi_t}{E_t} + \sigma \Xi_t, \quad (7)$$

where $\sigma \in [0, 1]$ is a weight factor, Ξ_t is the performance of the global model, i.e., test accuracy in the t -th round. Subsequently, we need to optimize the objective function under resource and task requirements, and the resource constraints essentially affect resource allocation strategies. Towards an exemplified scenario, which has limited CPU operating frequency, transmission power, and bandwidth, and should satisfy some time-related QoS requirements, we optimize the Cumulative Weighted Performance-Energy Metric (CWPEM) as:

$$\mathfrak{P}1 : \max_{\mathbb{S}_t, f_{t,n}, p_{t,n}, B_{t,n}} \sum_{t=1}^K r(\Xi_t, E_t) \quad (8a)$$

$$\text{s.t. } T_t \leq T_{\text{QoS}}, \quad (8b)$$

$$f_{t,n} \leq f_{n,\text{max}}, \quad (8c)$$

$$p_{t,n} \leq p_{n,\text{max}}, \quad (8d)$$

$$\sum_{n \in \mathbb{S}_t} B_{t,n} = B, \quad (8e)$$

where K is the total round of FL, T_{QoS} denote the QoS requirement for the FL competition time, B is the total bandwidth of the system, $f_{n,\text{max}}$ and $p_{n,\text{max}}$ denote the maximum local computation capacity and maximum transmission power of device n , respectively. Constraint (8b) ensures that the time consumption complies with the time-related QoS requirement. Constraints (8c) and (8d) specify that the allocated computational frequency and transmission power must not exceed the device's resource limitations. Finally, Constraint (8e) guarantees that the aggregate bandwidth utilized by the selected devices adheres to the system's bandwidth constraints.

A widely applied solver to the joint optimization problem $\mathfrak{P}1$ is DRL [10], [19], [20]. Generally, the dynamic process of the wireless FL can be formally modeled as the Markov Decision Process (MDP), denoted by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where \mathcal{S} is the state space with $\mathbf{s}_t \in \mathcal{S}$, and \mathbf{s}_t can be composed of two parts, i.e.,

$$\mathbf{s}_t = [\mathbf{s}_t^c, \mathbf{s}_t^m], \quad (9)$$

where \mathbf{s}_t^c represents parameters related to communication and computation, including $f_{t,\text{max}}$, $p_{t,\text{max}}$, $G_{t,n}$, $E_{t,n}^{\text{com}}$, and $E_{t,n}^{\text{comp}}$ in Section II-A2, while \mathbf{s}_t^m denotes the statistical parameters during training, which correspond to $|D_n|$, $F_n(\omega_n)$, and $\Xi(\mathbb{S}_{t-1})$ in Section II-A1. \mathcal{A} is the action space with $\mathbf{a}_t \in \mathcal{A}$ encompassing device selection and corresponding resource allocation results, i.e.,

$$\mathbf{a}_t = [\mathbb{S}_t, \mathbb{R}_t], \quad (10)$$

where \mathbb{R}_t represents the allocated resources, such as $f_{t,n}$, $p_{t,n}$, and $B_{t,n}$, to enable the full participation of devices within the subset \mathbb{S}_t , while we call such a subset as a *feasible* subset.¹ P is the transition function $P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$. R is the reward

¹Notably, if \mathbb{R}_t can only support the participation of a subset $\hat{\mathbb{S}}_t \subset \mathbb{S}_t$ of devices, we will denote the action as $[\hat{\mathbb{S}}_t, \mathbb{R}_t]$ to avoid the ambiguity while the feasible subset is denoted as $\hat{\mathbb{S}}_t$.

function calculated by (7) and γ is the discount factor. Generally, the NNs trained by DRL are only suitable for the current interactive environment, and thus, they need to be re-trained in a new environment. In addition, for different scenarios with distinctive resource constraints and task requirements, the NNs might be redesigned from scratch. For example, for two scenarios with different managed resources, the action space must be redesigned. Unfortunately, for a complicated problem that involves multiple interdependent decision-making tasks, the design and training of NNs becomes increasingly complex. Besides, since the number of neurons grows exponentially with the number of decision variables, such as f_n , p_n , and B_n per device, it commonly encounters scalability issues. Consequently, an excessively large decision-making action space adversely affects the efficacy of DRL. Therefore, we propose to introduce LLM to address these challenges.

III. DECOUPLING OF THE OPTIMIZATION PROBLEM

In this section, to make LLM compatible with the decision-making environment of wireless FL, we decouple the optimization problem. In particular, we decompose the joint optimization into two distinct sub-problems. Beforehand, we show that under a feasible device selection subset \mathbb{S}_t , there exists independence between the global model performance and the resource allocation.

Lemma 1: Given the feasible device selection action \mathbb{S}_t , the performance of global model Ξ is independent of the valid resource allocation action \mathbb{R}_t , where \mathbb{R}_t enable the full participation of devices within the subset \mathbb{S}_t , i.e.,

$$P(\Xi_t, \mathbb{R}_t | \mathbb{S}_t) = P(\Xi_t | \mathbb{S}_t) P(\mathbb{R}_t | \mathbb{S}_t). \quad (11)$$

Proof: By definition, all devices in \mathbb{S}_t are involved in the aggregation of FL, which implies no resource constraints in (8a) are violated. In this case, the parameters of the global model are completely determined by (4). Mathematically,

$$P(\Xi_t | \mathbb{R}_t, \mathbb{S}_t) = P(\Xi_t | \mathbb{S}_t). \quad (12)$$

Therefore, we have

$$\begin{aligned} P(\Xi_t, \mathbb{R}_t | \mathbb{S}_t) &= \frac{P(\Xi_t, \mathbb{R}_t, \mathbb{S}_t)}{P(\mathbb{S}_t)} \\ &= \frac{P(\mathbb{R}_t, \mathbb{S}_t) P(\Xi_t | \mathbb{S}_t)}{P(\mathbb{S}_t)} \\ &= P(\Xi_t | \mathbb{S}_t) P(\mathbb{R}_t | \mathbb{S}_t). \end{aligned} \quad (13)$$

We have the lemma. \blacksquare

Corollary 1: The performance of the global model Ξ is solely determined by the selected subset of participating devices \mathbb{S}_t , i.e.,

$$\Xi_t = \Xi(\mathbb{S}_t). \quad (14)$$

Afterward, we introduce some fundamental assumptions for the problem decoupling.

Assumption 1: For a general resource management problem with a feasible device subset \mathbb{S}_t , there always exists a function $g(\cdot)$ that can find the optimal solution of the joint optimization problem $\mathfrak{P}1$.

Assumption 2: For the joint optimization problem $\mathfrak{P}1$, given fixed and unchanging device resources (i.e.,

$f_{t,n} = f_n$, $p_{t,n} = p_n$ and $B_{t,n} = B_n$), an optimal feasible subset \mathbb{S}_t^* can always be found.

Based on these assumptions, we have the following theorem.

Theorem 1: Under Assumption 1 and Assumption 2, the joint optimization problem $\mathfrak{P}1$ can be equivalently decomposed into an energy-efficiency oriented resource management subproblem $\mathfrak{P}2$ and a device selection subproblem $\mathfrak{P}3$, where $\mathfrak{P}2$ and $\mathfrak{P}3$ are respectively defined as:

$$\begin{aligned} \mathfrak{P}2: \quad & \min_{f_{t,n}, p_{t,n}, b_{t,n}} \sum_{t=1}^K E_t \\ & \text{s.t. } \mathbb{S}_t \subset \mathcal{N}, \\ & \text{constraints (8b), (8c), (8d), (8e),} \end{aligned} \quad (15)$$

and

$$\begin{aligned} \mathfrak{P}3: \quad & \max_{\mathbb{S}_t} \sum_{t=1}^K r(\Xi(\mathbb{S}_t), E_t^*(\mathbb{S}_t)) \\ & \text{s.t. } T_{t,n} \leq T_{\text{QoS}}, \\ & f_{t,n}, p_{t,n}, B_{t,n} = g(\mathbb{S}_t). \end{aligned} \quad (16)$$

Proof: We first discuss the resource management issue in the problem $\mathfrak{P}1$. Assumption 1 implies that for a given feasible device subset \mathbb{S}_t , through $g(\cdot)$, the optimal resources can always be allocated to the device to minimize energy consumption, while meeting the required QoS and resource constraints. In other words, given any device subset \mathbb{S}_t , the minimum energy consumption can be calculated as $E_t^*(\mathbb{S}_t) = E_t(\mathbb{S}_t, g(\mathbb{S}_t))$ as in the problem $\mathfrak{P}2$.

Next, recalling the optimization objective $r(\Xi(\mathbb{S}_t), E_t)$ in the problem $\mathfrak{P}1$, it monotonically decreases with respect to E_t and increases with respect to $\Xi(\mathbb{S}_t)$, since Ξ is independent of resource allocation by Corollary 1. Based on Assumption 2, the optimal feasible subset \mathbb{S}_t^* can always be found for the sub-problem $\mathfrak{P}3$. Afterward, for any given device subset \mathbb{S}_t^* , it follows that

$$\begin{aligned} & \max_{\mathbb{S}_t, f_{t,n}, p_{t,n}, b_{t,n}} \sum_{t=1}^K r(\Xi(\mathbb{S}_t), E_t) \\ &= \max_{f_{t,n}, p_{t,n}, b_{t,n}} \sum_{t=1}^K r(\Xi(\mathbb{S}_t^*), E_t) \\ &= \sum_{t=1}^K r(\Xi(\mathbb{S}_t^*), E_t^*(\mathbb{S}_t^*)). \end{aligned} \quad (17)$$

Therefore, $\mathfrak{P}1$ is equivalent to solving $\mathfrak{P}2$ and $\mathfrak{P}3$ separately, and the optimal solution is given by $(\mathbb{S}_t^*, g(\mathbb{S}_t^*))$. \blacksquare

Remark: Theorem 1 provides the formal basis for decoupling the joint optimization problem. This allows us to deal with the two sub-problems separately and reduces computational complexity by searching over a smaller-dimensional solution space. Accordingly, we can introduce LLM to solve $\mathfrak{P}3$ conveniently, while resorting to a tool-aided solution of $\mathfrak{P}2$. Notably, both sub-problems still involve a calibrated joint design for decision making: For the subproblem $\mathfrak{P}2$, the function $g(\cdot)$ satisfying Assumption 1 can be practically guaranteed by optimization tools. For $\mathfrak{P}2$, the policy itself is required to make optimal device selection decisions under the condition of the resource allocation mechanism $g(\cdot)$.

IV. TOOL-AIDED EVOLUTIONARY LLM SCHEME

In this section, we first introduce the tool-aided LLM towards generating an effective device selection and resource

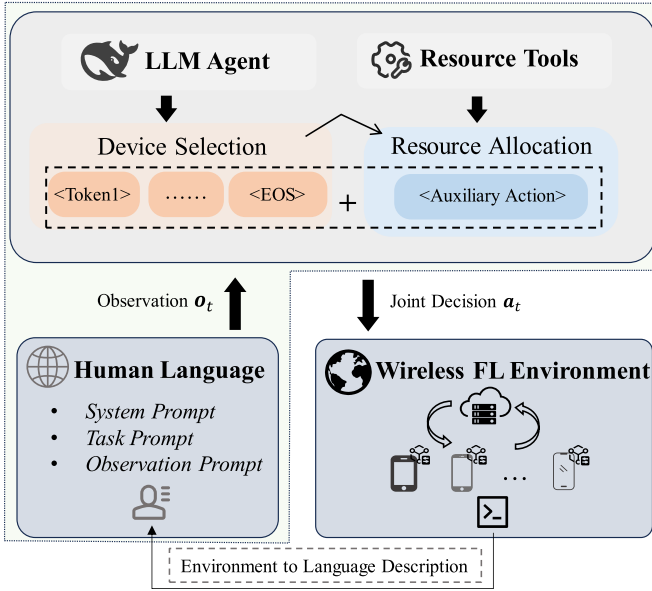


Fig. 2. An illustration of the inference process by the tool-aided LLM.

allocation policy. Afterward, on top of convex optimization-based resource management tools, we introduce a model-based, tool-assisted training framework. Finally, we discuss how to train the LLM by GRPO, thus improving its decision-making capabilities.

A. Tool-Aided LLM for Efficient Policy Generation

The decision-making pipeline of the proposed tool-aided LLM is shown in Fig. 2, which is motivated by the mathematical findings in Section III and can be divided into two parts: a generalizable LLM to provide device section results for $\mathfrak{P}3$ and an optimization tool to yield scenario-specific contextual resource allocation outcomes for $\mathfrak{P}2$. To generate a resource-efficient policy, at each communication round t , the current observation of the FL environment will be captured and described directly in language. Correspondingly, the state space \mathcal{S} in the MDP \mathcal{M} is converted to text-based observation \mathcal{O} with $\mathbf{o}_t \in \mathcal{O}$, and

$$\mathbf{o}_t = [\mathbf{s}_t^c, \mathbf{s}_t^m, \mathbf{s}_t^{\text{text}}], \quad (18)$$

where $\mathbf{s}_t^{\text{text}}$ represents the added natural language description. This linguistic representation \mathbf{o}_t , which will be elaborated later, serves as the input to the LLM. Consequently, it efficiently adapts to diverse state spaces across various scenarios without requiring additional tuning of NN architectures. Let \mathcal{V} represent the LLM's entire token space. For a given observation \mathbf{o}_t , the LLM first computes a vector of logits over all tokens in \mathcal{V} , i.e., $L_v(\mathbf{o}_t)$. To ensure that the generated action corresponds to a valid device, we define a subset of this token space $\mathcal{I} \subset \mathcal{V}$, as the set of valid action tokens, where each token represents a unique device index $n \in \mathcal{N}$. Therefore, the logits over all valid tokens can be expressed as:

$$L_v^I(\mathbf{o}_t) = \begin{cases} L_v(\mathbf{o}_t) & \text{if } v \in \mathcal{I} \\ -\infty & \text{if } v \notin \mathcal{I}. \end{cases} \quad (19)$$

A softmax function is then applied to this masked logit vector to produce the final probability distribution over the vocabulary, that is, $\pi_\theta(\cdot|\mathbf{o}_t) := \pi_\theta(v|\mathbf{o}_t) = \frac{\exp(L_v^I(\mathbf{o}_t))}{\sum_{u \in \mathcal{V}} \exp(L_u^I(\mathbf{o}_t))}$. Then, the LLM generates a text response $\tilde{\mathbf{a}}_t$ based on the current language description according to its strategy π_θ , i.e.,

$$\tilde{\mathbf{a}}_t \sim \pi_\theta(\cdot|\mathbf{o}_t), \quad (20)$$

which is interpreted as the decision action for device selection \mathbb{S}_t . Upon selecting the subset \mathbb{S}_t of devices, the LLM invokes an external resource *tool* to derive the resource allocation action \mathbb{R}_t for the selected devices. Finally, the joint decisions $\mathbf{a}_t = [\mathbb{S}_t, \mathbb{R}_t]$ are applied to the wireless FL environment. Thus, the next observation can be obtained by the environment, i.e.,

$$P(\mathbf{o}_{t+1}|\mathbf{o}_t, \mathbf{a}_t). \quad (21)$$

Afterward, the feedback provided to the LLM-based policy generator triggers the next round of efficiency optimization.

1) *The LLM for $\mathfrak{P}3$* : As shown in Fig. 3, the LLM takes the linguistic description \mathbf{o}_t of the environmental observations as input and generates the device selection decisions \mathbb{S}_t , corresponding to the indices of devices. Prominently, the prompts, which include system prompt, task prompt, and observation prompt, are specially tailored for the wireless FL problem $\mathfrak{P}1$. Notably, system and task prompts configure the LLM as an agent for FL. These prompts guide the LLM in analyzing problems based on the provided task description, thereby increasing the likelihood of generating task-related token outputs. Besides, the observation prompt integrates all the state parameters \mathbf{s}_t^c and \mathbf{s}_t^m required for the current task, furnishing the LLM with a comprehensive description of the current environment for decision-making. Additionally, to ensure compliance with operational constraints, the prompt also specifies the desired output format of \mathbb{S}_t as a token-based response $\tilde{\mathbf{a}}_t$. On this basis, we fine-tune the LLM to improve the performance.

2) *The Tool for $\mathfrak{P}2$* : For the tool-aided LLM, the tool shall have the capability to compute optimal resource allocation solutions to $\mathfrak{P}2$ under a unique scenario, so that LLM can be compatible with the specific tool to make efficient joint decisions in response to this specific scenario. Numerous qualified algorithms [19], [20] can be leveraged. In this work, due to its implementation simplicity, we adopt the alternative direction algorithm ALTD proposed by [20]. Notably, for $\mathfrak{P}2$, based on device selection \mathbb{S}_t , ALTD optimizes allocated resources \mathbb{R}_t including the CPU frequency $f_{t,n}$, transmission power $p_{t,n}$ and bandwidth allocation $B_{t,n}$ of networking devices under the constraints of average bandwidth allocation and QoS requirements, thus allowing the participating devices to perform FL with the lowest energy consumption.

B. Model-Based Virtual Environment

Due to the significant communication overhead associated with online interactions in wireless FL, this section proposes a virtual environment based on an offline model-based approach and tool-assisted computations. Benefiting from the following lemma, the learning of the virtual environment can avoid the extensive use of huge end-to-end expert data.

Question

System Prompt: As a federated training agent, you are responsible for selecting the most suitable devices from the device pool to optimize performance in the current round.

Task Prompt: Select 4 indexes from given 20 devices to participate in federated learning training based on the return information in last round and the devices' states. Please give 4 indexes as a result directly without additional information.

Observation Prompt: For last round, the select devices are [1, 6, 4, 8], and the time consumption is 15.0, the energy consumption is 3.8031, the accuracy of global model is 0.4372. The total bandwidth of the system is $2e7$. The reference QoS time is 15.00s. The current communication is round 2. The states information of each device is shown in the following table:

device index	max power	max frequency	channel gain	number of compute cycles	data size	local loss	local loss after trained	inner-product between local model and global	percentage of same sign between local model and global model	last selected round	selected times
0	0.829053	2049452125.9596	4.8439e-07	700000	2232	1874.7955	\	\	\	\	\
1	0.564018	1229342133.7070	1.1921e-06	700000	1527	2114.6651	6401.0288	0.3673	0.4415	1	1
.....											
18	0.588518	3479022664.3894	1.4571e-06	700000	2970	2844.2611	\	\	\	\	\
19	0.014806	1573110420.5726	4.9860e-07	700000	2668	6298.7290	\	\	\	\	\

Output Limitation: Please select 4 indexes from given 20 device indexes as a result. Give one result directly. Your response consists only of indexes.

FL Agent Response

'5' '2' '4' '18'

Fig. 3. An example of input and output for device selection decision-making by LLM.

Lemma 2: The state transition of the wireless FL environment can be decomposed into two parts:

$$P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) = P(\mathbf{s}_{t+1}^c|\mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))P(\mathbf{s}_{t+1}^m|\mathbf{s}_t^m, \mathbb{S}_t). \quad (22)$$

Proof: The state transition of the wireless RL environment can be represented as:

$$\begin{aligned} P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) &= P(\mathbf{s}_{t+1}^m, \mathbf{s}_{t+1}^c|\mathbf{s}_t, \mathbf{a}_t) \\ &= P(\mathbf{s}_{t+1}^m, \mathbf{s}_{t+1}^c|(\mathbf{s}_t^m, \mathbf{s}_t^c), (\mathbb{S}_t, \mathbb{R}_t)). \end{aligned} \quad (23)$$

By definition, \mathbf{s}^c and \mathbf{s}^m are conditional independent. Therefore, we can have

$$\begin{aligned} P(\mathbf{s}_{t+1}^m, \mathbf{s}_{t+1}^c|(\mathbf{s}_t^m, \mathbf{s}_t^c), (\mathbb{S}_t, \mathbb{R}_t)) \\ = P(\mathbf{s}_{t+1}^m|(\mathbf{s}_t^m, \mathbf{s}_t^c), (\mathbb{S}_t, \mathbb{R}_t))P(\mathbf{s}_{t+1}^c|(\mathbf{s}_t^m, \mathbf{s}_t^c), (\mathbb{S}_t, \mathbb{R}_t)). \end{aligned} \quad (24)$$

Similarly, according to the independence of \mathbf{s}_{t+1}^m and \mathbf{s}_t^c , as well as \mathbf{s}_{t+1}^c and \mathbf{s}_t^m , we have

$$P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) = P(\mathbf{s}_{t+1}^m|\mathbf{s}_t^m, (\mathbb{S}_t, \mathbb{R}_t))P(\mathbf{s}_{t+1}^c|\mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t)). \quad (25)$$

Recalling that by definition, \mathbf{s}^m is $\{|D_n|, F_n(\omega_n), \Xi(\mathbb{S}_{t-1})\}$ -related. Then, by Lemma 1, we have the conclusion. ■

Lemma 2 implies that the wireless FL environment can be simulated by a *system part* and an *FL statistics part*, which is

illustrated in Fig. 4. For the system part, it only focuses on the communication and computation status of devices according to the given resources, i.e., $P(\mathbf{s}_{t+1}^c|\mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))$. Therefore, the observation and consumption can be recorded and simulated by the resource tool. Specifically, based on \mathbb{R}_t and \mathbb{S}_t , we can leverage the tool to calculate the consumed time T_t and energy E_t , and update the state \mathbf{s}_{t+1}^c .

However, for characterizing the statistics part $P(\mathbf{s}_{t+1}^m|\mathbf{s}_t^m, \mathbb{S}_t)$, it remains challenging when dynamically selected devices participate in training and aggregation. To address this, we employ a model to approximate the convergence behavior and performance of the global model during FL training. In other words, simulating the statistics part $P(\mathbf{s}_{t+1}^m|\mathbf{s}_t^m, \mathbb{S}_t)$ means to predict \mathbf{s}_{t+1}^m including $\Xi(\mathbb{S}_t)$ based on the current observation \mathbf{s}_t^m , device selection action \mathbb{S}_t and the historical accuracy $\Xi(\mathbb{S}_{t-1})$. Generally, \mathbf{s}_t^m can be classified as the state information pertinent to the selected devices $\mathbf{s}_t^m(s)$ and that of all devices $\mathbf{s}_t^m(l)$. In particular, $\mathbf{s}_t^m(s)$ includes post-training local loss $F_n(\omega_n)$, the inner-product $\langle \nabla F_n(\omega_n^t), \nabla F_G(\omega_G^t) \rangle$, and the percentage of the same sign $e(\omega_n^t, \omega_G^t)$ shared between local and global models (i.e., ω_n^t and ω_G^t) for all selected devices $n \in \mathbb{S}_t$. Meanwhile, $\mathbf{s}_t^m(l)$ consists of the local dataset size $|D_n|$ (which remains constant during training) and locally computed loss l_n of other non-selected devices. We denote

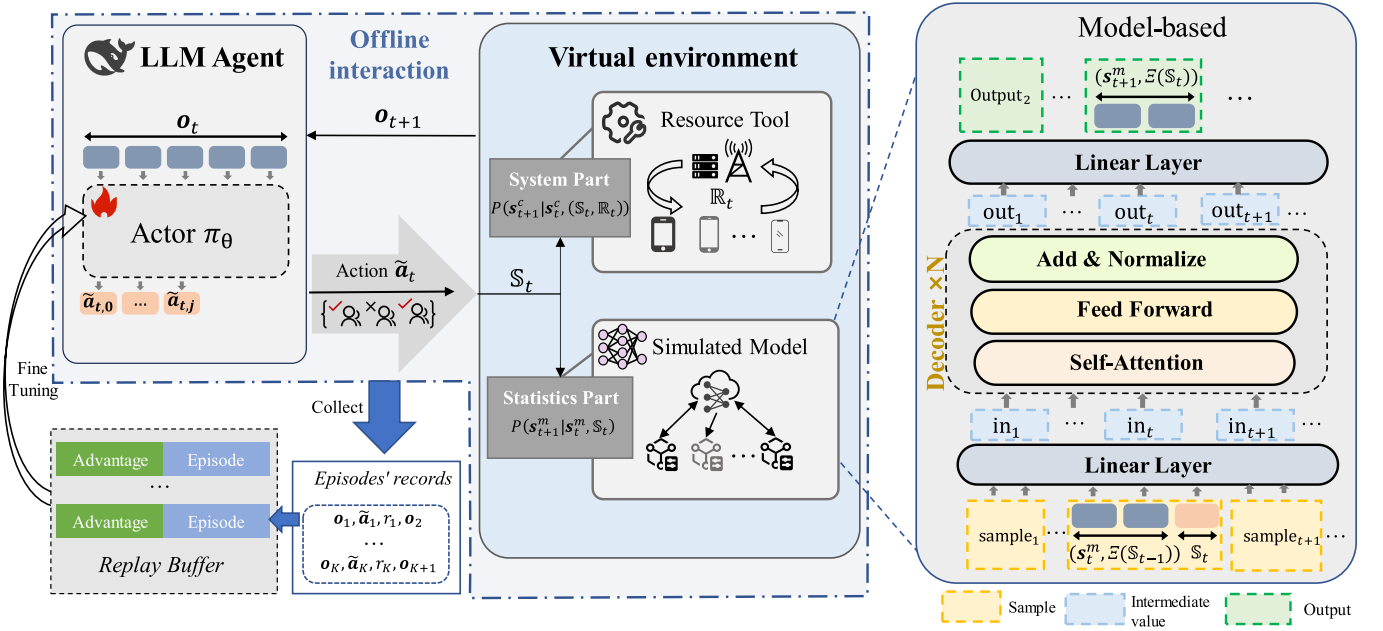


Fig. 4. Illustration of the evolutionary virtual environment and GRPO-based training in T-ELLM.

the tuple $((s_t^m, \Xi(S_{t-1}), S_t), (s_{t+1}^m, \Xi(S_t)))$ as one sample of the environment model, while several continuous samples constitute a trajectory. The NN architecture of the simulated model is shown in the right part of Fig. 4, while the details of the training process are shown in Algorithm 1. During training, we optimize the model with a Mean Squared Error (MSE) loss.

Generally, the reward r_t can be computed as in (7), except when the resource tool finds no solution exists for $\mathfrak{P}2$, the action S_t will be given a penalty, i.e., $r_t = 0$. Then, alongside the reward r_t , the virtual environment can respond to $(s_t^m, \Xi(S_{t-1}), S_t)$, and accurately yield the next state (s_{t+1}^c, s_{t+1}^m) .

C. GRPO-Based LLM Training

The aforementioned virtual environment lays the very foundation for updating the LLM agent in an RL manner. Specifically, we employ GRPO [38] to enhance the inference capability of LLM. As a variant of Proximal Policy Optimization (PPO) [41], GRPO also adopts twin policy models, where an old policy $\pi_{\theta_{\text{old}}}$ is used for sampling actions $\tilde{\mathbf{a}}_t$ and accumulating transitional records $(\mathbf{o}_t, \tilde{\mathbf{a}}_t, r_t, \mathbf{o}_{t+1})$ in a replay buffer. Here, r_t is the reward in the t -th round calculated as in Section IV-B. All samples are described in language, a specific example as shown in Fig. 3. For a batch of K -length episodes' records, where each episode i is denoted as $\{r_1, \dots, r_t, \dots, r_K\}_i$, we normalize these rewards with the t -wise average and standard deviation across episodes within the batch, i.e., $\{\tilde{r}_t\}_i = \frac{\{r_t\}_i - \text{mean}}{\text{std}}$. Then the t -th advantage in i -th episode can be calculated as:

$$\{A_t\}_i = \sum_{j \geq t} \{\tilde{r}_j\}_i. \quad (26)$$

Based on the $\pi_{\theta_{\text{old}}}$ -induced records in the relay buffer, it becomes ready to update the current policy model π_{θ} . Notably,

due to the autoregressive characteristics of language tasks, GRPO [38] can regressively update LLMs by taking a single token as an action. While in a wireless FL environment, an effective action is the entire response rather than a token. Therefore, wireless FL is completely different from natural language processing tasks. Let $\mathbf{a}_{t,j}$ represent the action taken in round t and j -th token, thus $\mathbf{o}_{t,<j} = [\mathbf{o}_{t,<j-1}, \tilde{\mathbf{a}}_{t,j}]$, and $\mathbf{o}_{t,<0} = \mathbf{o}_t$. In this case, the conditional probability of each action $\tilde{\mathbf{a}}_t$ can be calculated as:

$$\pi_{\theta}(\tilde{\mathbf{a}}_t | \mathbf{o}_t) = \prod_j \pi_{\theta}(\tilde{\mathbf{a}}_{t,j} | \mathbf{o}_{t,<j-1}). \quad (27)$$

Next, we adopt GRPO to optimize the LLM by maximizing

$$J_{\text{GRPO}}(\theta) = \mathbb{E} \left[\min \left(\frac{\pi_{\theta}(\tilde{\mathbf{a}}_t | \mathbf{o}_t)}{\pi_{\theta_{\text{old}}}(\tilde{\mathbf{a}}_t | \mathbf{o}_t)} A_t, \text{clip} \left(\frac{\pi_{\theta}(\tilde{\mathbf{a}}_t | \mathbf{o}_t)}{\pi_{\theta_{\text{old}}}(\tilde{\mathbf{a}}_t | \mathbf{o}_t)}, 1 - \epsilon, 1 + \epsilon \right) A_t \right) \right], \quad (28)$$

where ϵ is a clipping-related hyperparameter for stabilizing training.

Finally, the implementation process of the T-ELLM scheme is illustrated in Algorithm 1.

V. THEORETICAL ACCURACY OF THE POLICY LEARNED FROM THE VIRTUAL ENVIRONMENT

This section theoretically validates the proposed model-based virtual environment by establishing a generalized simulation lemma. Specifically, we demonstrate that the overall discrepancy between the simulated environment and the real environment is provably bounded, i.e., the convergence properties of FL can be faithfully replicated under the model-based environment. Therefore, we further show that the advantage function, i.e., cumulative normalized reward under the learned virtual environment in (26), remains within a

Algorithm 1 GRPO-Based T-ELLM Scheme in the Model-Based Virtual Environment

Training for model-based virtual environment

 Collect T_m trajectories $[(\mathbf{s}_t^m, \Xi(\mathbb{S}_{t-1}), \mathbb{S}_t), (\mathbf{s}_{t+1}^m, \Xi(\mathbb{S}_t))]$ as the training dataset.

- 1: Add two linear layers (i.e., `linear1` and `linear2`) for embedding to GPT-2 model and initialize all parameters ϕ .
- 2: **for** each epoch **do**
- 3: **for** each batch **do**
- 4: **parfor** all samples in the same trajectory **do**
- 5: $\text{in}_t = \text{linear1}(\mathbf{s}_t^m, \Xi(\mathbb{S}_{t-1}), \mathbb{S}_t)$.
- 6: $\text{out}_t = \text{GPT}(\text{in}_1, \dots, \text{in}_{t-1})$.
- 7: $(\text{blue}\hat{\mathbf{s}}_{t+1}^m, \hat{\Xi}(\mathbb{S}_t)) = \text{linear2}(\text{out}_t)$.
- 8: $\text{loss}_t = \text{MSE}((\hat{\mathbf{s}}_{t+1}^m, \hat{\Xi}(\mathbb{S}_t)), (\mathbf{s}_{t+1}^m, \Xi(\mathbb{S}_t)))$.
- 9: **end parfor**
- 10: $\text{loss}_{\text{traj}} = \text{mean}(\text{loss}_t)$.
- 11: **end for**
- 12: $\text{loss}_{\text{batch}} = \text{mean}(\text{loss}_{\text{traj}})$.
- 13: $\text{loss}_{\text{batch}}$ backward.
- 14: update ϕ and two linear layers.

end for
Fine-tuning for tool-aided evolutionary LLM

 Initialize all the parameters: policy model π_θ , hyperparameters ϵ, λ , the trained simulated model.

- 1: Initialize $\pi_{\theta_{\text{old}}} = \pi_\theta$.
 - 2: **for** each iteration **do**
 - 3: $\pi_{\theta_{\text{old}}} = \pi_\theta$.
 - 4: Episode set $\text{Epi} = \emptyset$.
 - 5: **for** $i = 1, \dots, \text{batch}$ **do**
 - 6: $\text{Epi}_i = \emptyset$
 - 7: **for** communication round $t = 1, \dots, K$ **do**
 - 8: Obtain the observation \mathbf{o}_t .
 - 9: Sample $\mathbf{a}_t \sim \pi_{\theta_{\text{old}}}(\cdot | \mathbf{o}_t)$.
 - 10: Execute \mathbf{a}_t , receive r_t .
 - 11: Get \mathbf{o}_{t+1} via resource tool and simulated model.
 - 12: $\text{record}_t = (\mathbf{o}_t, \tilde{\mathbf{a}}_t, r_t, \mathbf{o}_{t+1})$
 - 13: Store $\text{Epi}_i \leftarrow \text{Epi}_i \cup \text{record}_t$.
 - 14: **end for**
 - 15: $\text{Epi} \leftarrow \text{Epi} \cup \text{Epi}_i$.
 - 16: **end for**
 - 17: Calculate all advantages A_t based on (26).
 - 18: **for** each update step **do**
 - 19: Update π_θ by maximizing (28) with Epi .
 - 20: **end for**
 - 21: **end for**
-

bounded deviation from the real environment, ensuring similar convergence behavior of LLMs in both scenarios.

The MDP of the entire wireless FL system is expressed as $\mathcal{M} = (\mathcal{O}, \mathcal{A}, P, \mathcal{R}, \gamma)$. Specially, the state space is text-based observation $\mathbf{o}_t = [\mathbf{s}_t^c, \mathbf{s}_t^m, \mathbf{s}_t^{\text{ext}}]$. And the added natural language description $\mathbf{s}_t^{\text{ext}}$ as the prompt for LLM is independent of the wireless FL. Similarly, let $\hat{\mathcal{M}}$ denote the MDP representing the complete model-based virtual environment, and \hat{P} its corresponding transition function. To formally characterize the difference between the real environment \mathcal{M}

and the model-based approximation $\hat{\mathcal{M}}$, we first establish the following assumptions and define the return-based advantage function used in our analysis.

Assumption 3: Let \mathcal{M} and $\hat{\mathcal{M}}$ be two MDPs defined over the same state and action spaces. Suppose that the normalized reward functions satisfy the following uniform deviation bound for all state-action pairs $(\mathbf{o}_t, \mathbf{a}_t)$, the reward functions satisfy

$$|\tilde{r}_{\hat{\mathcal{M}}}(\mathbf{o}_t, \mathbf{a}_t) - \tilde{r}_{\mathcal{M}}(\mathbf{o}_t, \mathbf{a}_t)| < \tilde{R}_{\max}, \quad (29)$$

where \tilde{R}_{\max} denotes a constant.

Definition 1: For any policy π and state \mathbf{o}_t , consistent with (26), define the return-based advantage function as:

$$A_P^\pi(\mathbf{o}, \mathbf{a}) := \mathbb{E}_{P, \pi} \left[\sum_{j=t}^{K-1} \tilde{r}(\mathbf{o}_j, \mathbf{a}_j) \right], \quad (30)$$

where $(\mathbf{o}_t, \mathbf{a}_t) = (\mathbf{o}, \mathbf{a})$, $\tilde{r}(\mathbf{o}, \mathbf{a})$ denotes a normalized reward function satisfying $|\tilde{r}(\mathbf{o}, \mathbf{a})| \leq \tilde{R}_{\max}$.

Our approach separately characterizes the system and statistics parts instead of using a real wireless FL environment. The system part uses resource tools to simulate communication and computing costs, where fixed resource (e.g., bandwidth and transmission power) settings often lead to deterministic state transitions $P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))$, due to the explicit mathematical expressions in Section II-A2. On the other hand, the MDP corresponding to the statistical component is denoted as $\mathcal{M}_d = (\mathcal{S}^m, \mathcal{A}^m, P^m, \mathcal{R}^m, \gamma)$, where, as described in Section IV, the state is defined as $\mathbf{s}_t^m = \{\mathbf{s}_t^m(s), \mathbf{s}_t^m(l)\} \in \mathcal{S}^m$. The decision action corresponds to selecting a subset \mathbb{S}_t of devices. The transition probability is given by $P^m = P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t)$, and the reward is defined as the accuracy of the global model, denoted by $\Xi(\mathbb{S}_t)$. To approximate this statistical MDP, we employ a trained model $\hat{\mathcal{M}}_d$, whose transition dynamics are captured by the learned function \hat{P}^m .

Assumption 4 (Bounded Transition Discrepancy in TV Distance): Let P^m and \hat{P}^m denote the transition functions of \mathcal{M}_d and $\hat{\mathcal{M}}_d$, respectively. Suppose the transition model discrepancy is bounded in Total Variation (TV) as:

$$\sup_{\mathbf{s}_t^m, \mathbb{S}_t} D_{\text{TV}} \left(\hat{P}(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t), P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t) \right) \leq \epsilon, \quad (31)$$

where $D_{\text{TV}}(p, q)$ denotes the TV [42] distance between probability distributions p and q .

Lemma 3: Under Assumption 4, let P and \hat{P} denote the transition functions of \mathcal{M} and $\hat{\mathcal{M}}$, respectively. The transition model discrepancy is bounded in total variation as

$$\sup_{\mathbf{o}_t, \mathbf{a}_t} \left\| \hat{P}(\mathbf{o}_{t+1} | \mathbf{o}_t, \mathbf{a}_t) - P(\mathbf{o}_{t+1} | \mathbf{o}_t, \mathbf{a}_t) \right\|_1 \leq 2\epsilon. \quad (32)$$

Proof: By definition, the transition probability function P of the full environment model can be decomposed as:

$$P(\mathbf{o}_{t+1} | \mathbf{o}_t, \mathbf{a}_t) = P(\mathbf{s}_{t+1}^m, \mathbf{s}_{t+1}^c, \mathbf{s}_{t+1}^{\text{ext}} | (\mathbf{s}_t^m, \mathbf{s}_t^c, \mathbf{s}_t^{\text{ext}}), (\mathbb{S}_t, \mathbb{R}_t)),$$

where $\mathbf{s}_t^{\text{ext}}$ denotes the language inputs to the LLM, which is independent of other variables and can thus be treated as a constant. Accordingly, consistent with Lemma 2, the transition

function P can be simplified without loss of theoretical generality as follows:

$$\begin{aligned} P(\mathbf{o}_{t+1}|\mathbf{o}_t, \mathbf{a}_t) &= P(\mathbf{s}_{t+1}^m, \mathbf{s}_{t+1}^c | (\mathbf{s}_t^m, \mathbf{s}_t^c)(\mathbb{S}_t, \mathbb{R}_t)) \\ &= P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t), \end{aligned} \quad (33)$$

where $P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))$ is the transition function of system part, which is determined by the resource tool. Similarly, the transition function \hat{P} in the virtual environment can be obtained as:

$$\hat{P}(\mathbf{o}_{t+1}|\mathbf{o}_t, \mathbf{a}_t) = P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t))\hat{P}(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t).$$

Therefore, under Assumption 4, we have

$$\begin{aligned} &\left\| \hat{P}(\mathbf{o}_{t+1}|\mathbf{o}_t, \mathbf{a}_t) - P(\mathbf{o}_{t+1}|\mathbf{o}_t, \mathbf{a}_t) \right\|_1 \\ &= \sum_{\mathbf{s}_{t+1}^c, \mathbf{s}_{t+1}^m} \left| P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t)) \cdot \hat{P}(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t) \right. \\ &\quad \left. - P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t)) \cdot P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t) \right| \\ &= \sum_{\mathbf{s}_{t+1}^c} P(\mathbf{s}_{t+1}^c | \mathbf{s}_t^c, (\mathbb{S}_t, \mathbb{R}_t)) \cdot \sum_{\mathbf{s}_{t+1}^m} \left| \hat{P}(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t) \right. \\ &\quad \left. - P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t) \right| \\ &= \sum_{\mathbf{s}_{t+1}^c} P(\mathbf{s}_{t+1}^c | \cdot) \cdot \left\| \hat{P}(\cdot | \mathbf{s}_t^m, \mathbb{S}_t) - P(\cdot | \mathbf{s}_t^m, \mathbb{S}_t) \right\|_1 \\ &\stackrel{(a)}{=} \left\| \hat{P}(\cdot | \mathbf{s}_t^m, \mathbb{S}_t) - P(\cdot | \mathbf{s}_t^m, \mathbb{S}_t) \right\|_1 \\ &\stackrel{(b)}{=} 2D_{\text{TV}}\left(\hat{P}(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t), P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m, \mathbb{S}_t)\right) \stackrel{(c)}{\leq} 2\varepsilon. \end{aligned} \quad (34)$$

Specifically, (a) results from isolating the system transition probability as a common factor. (b) is derived by converting the ℓ_1 -norm into TV distance via its standard definition $D_{\text{TV}}(p, q) = \frac{1}{2}\|p - q\|_1$. (c) applies Assumption 4 that the TV between the model and true statistical transitions is at most ε . Hence, we have the corollary. ■

The following theorem characterizes the bounded deviation of the return-based advantage under model-based dynamics.

Theorem 2 (Generalized Simulation Lemma for Return-Based Advantage Functions): Under Assumption 3 and Assumption 4, for any fixed policy π , the difference in cumulative normalized rewards over a finite horizon K between the model-based MDP $\tilde{\mathcal{M}}$ and the true environment \mathcal{M} is bounded by

$$\left| A_{\hat{P}}^\pi(\mathbf{o}, \mathbf{a}) - A_P^\pi(\mathbf{o}, \mathbf{a}) \right| \leq (K^2 - K)\tilde{R}_{\max} \cdot \varepsilon. \quad (35)$$

Proof: By definition,

$$\begin{aligned} &\left| A_{\hat{P}}^\pi(\mathbf{o}, \mathbf{a}) - A_P^\pi(\mathbf{o}, \mathbf{a}) \right| \\ &= \left| \mathbb{E}_{\hat{P}, \pi} \left[\sum_{j=t}^{K-1} \tilde{r}(\mathbf{o}_j, \mathbf{a}_j) \right] - \mathbb{E}_{P, \pi} \left[\sum_{j=t}^{K-1} \tilde{r}(\mathbf{o}_j, \mathbf{a}_j) \right] \right| \\ &\leq \sum_{t=0}^{K-1} \left| \mathbb{E}_{\hat{P}, \pi} [\tilde{r}_t] - \mathbb{E}_{P, \pi} [\tilde{r}_t] \right|. \end{aligned} \quad (36)$$

Since $\tilde{r}(\mathbf{o}_j, \mathbf{a}_j)$ is bounded by \tilde{R}_{\max} according to Assumption 3, and by TV Inequality, shown in Prop. 4.2 of [43], we have

$$\left| \mathbb{E}_{\hat{P}, \pi} [\tilde{r}_t] - \mathbb{E}_{P, \pi} [\tilde{r}_t] \right| \leq \tilde{R}_{\max} \cdot \left\| \hat{P}_t^\pi - P_t^\pi \right\|_1, \quad (37)$$

where P_t^π and \hat{P}_t^π denote the distributions over t -step state-action trajectories induced by policy π under transition models P and \hat{P} , respectively. The total variation between K -step distributions grows at most linearly with the horizon length, as the distributional shift at step t reflects the accumulation of per-step transition errors over the entire trajectory up to time t [44]. In other words, by Lemma 3,

$$\left\| \hat{P}_t^\pi - P_t^\pi \right\|_1 \leq 2t \cdot \varepsilon. \quad (38)$$

By merging (37) and (38) into (36), we have the theorem after simple mathematical manipulations. ■

Remark: Theorem 2 result shows that the cumulative advantage error is controllable, growing polynomially with horizon K and linearly with the reward bound and transition error ε . Therefore, there are several practical implications. First, the framework is most effective for scenarios with a finite horizon, as the cumulative error may become significant in very long-running tasks. Second, the performance bound is directly tied to the fidelity of the virtual environment. Finally, the reliance on \tilde{R}_{\max} underscores the importance of a well-designed and stable reward function, as volatility or poor scaling can loosen the bound. Fortunately, with normalized rewards ($R_{\max} \leq 1$) and a Transformer-based transition model, ε can typically be reduced to a sufficiently accurate level. Together with proper control of K , the deviation between model-based and true advantages remains tightly bound.

VI. SIMULATION RESULTS

A. Experimental Setup

We consider an image classification task in a wireless FL scenario, as discussed in Section II, where the total number of devices is $N = 20$. In the experiments, we use the well-known MNIST dataset for FL training of a Convolutional Neural Network (CNN) model with the cross-entropy loss function. To simulate data heterogeneity across devices, we sample label ratios and dataset sizes from a Dirichlet distribution parameterized by α , which controls the degree of non-IIDness. Notably, a smaller α leads to more non-IID data, while a larger α results in more homogeneous data. We set $\alpha = 0.2$ as the default to simulate the non-IID data distribution. Other default configurations of simulation parameters are specified in the Table III.

For the training of T-ELLM, we employ the ALTD [20] as a complementary tool for dynamic management of CPU frequency and transmission power, as well as equal, fixed bandwidth allocation. In addition, the environmental model, which is adapted from the decoder-only part of the GPT-2-small architecture [26], [45] with two extra linear layers, is trained first. Subsequently, a LLaMA-3-1B-based [28] T-ELLM is used for GRPO [38]-based policy learning and generation. Besides, to evaluate the performance of the proposed T-ELLM, we compare it with the following baselines.

- FedAvg Tool [2]: A specific proportion of clients are randomly selected to participate in each round of FL training. Since it does not have the function of resource

TABLE I
A SUMMARY OF DIFFERENCES WITH RELATED LITERATURE ON EFFICIENT FL DECISION-MAKING

	Performance Guarantee	Adaptability to Environments	Artificial Design Complexity	Interaction Cost	Brief Description
Heuristics-based [7], [8]	●	●	●	●	Reliance on expert hand-designed utility functions
Optimization-based [18]–[20], [22]	●	●	○	●	Mainly applicable for optimization problems with expressions
RL-based [9]–[13], [23]–[25]	●	○	●	○	Environment-specific interactions with calibrated states
T-ELLM	●	●	●	●	Evolution in virtual environments through GRPO with tool assistance

Notations: ● Excellent; ● Medium; ○ Poor.

TABLE II
LIST OF KEY NOTATIONS USED IN THE PAPER

Notion	Definition
t	Index of the rounds
K	The total number of rounds
n	Index of device
N	Total number of devices
D_n	Local dataset of device n
D	Entire dataset
$ D $	The size of D
ω_n	Weight parameter of local model of device n
ω_G	Weight parameter of global model
\mathbb{S}_t	The set of indexes of the selected devices in round t
$\Xi(\mathbb{S}_t)$	Test accuracy of global model in round t
I	The number of iterations for local training
C	The number of CPU cycles to process one sample
ζ	The effective switch capacitance constant
$f_{t,n}$	CPU frequency of device n during the t -th round
$f_{n,\max}$	The maximum CPU frequency of device n
$p_{t,n}$	Transmission power of device n during the t -th round
$p_{n,\max}$	The maximum transmission power of device n
$B_{t,n}$	Bandwidth of device n during the t -th round
B	The total bandwidth of the system
$G_{t,n}$	Channel coefficient of device n during the t -th round
N_0	The noise power spectral density
s_n	The size (in bits) of the model parameters transmitted by device n
$T_{t,n}^{\text{cmp}}$	The time for local training at device n during the t -th round
$E_{t,n}^{\text{cmp}}$	The energy for local training at device n during the t -th round
$T_{t,n}^{\text{com}}$	The time for data uploading at device n during the t -th round
$E_{t,n}^{\text{com}}$	The energy for data uploading at device n during the t -th round
$v_{t,n}$	The transmission rate of device n during the t -th round
E_t	The total energy consumption during the t -th round
T_t	The total time consumption during the t -th round
\mathbf{s}_t^m	The statistical parameters during FL training
\mathbf{s}_t^c	State parameters related to communication and computation
\mathbb{R}_t	The allocated resources
\mathbf{a}_t	Joint action of device selection and resource allocation
\mathbf{o}_t	Linguistic representations of environmental state
$\tilde{\mathbf{a}}_t$	Text-based decision action for device selection

allocation itself, we use the ALTD tool to make resource allocation decisions.

- ϵ -Greedy Tool [7]: A variant of the Oort algorithm. Specifically, energy considerations are added, and tools are leveraged to obtain time and energy reference. Moreover, we assume all the device time and energy

TABLE III
DEFAULT VALUES OF SIMULATION PARAMETERS

Parameters	Value
Total number of devices N	20
Number of rounds K (Episodic length)	100
Number of local iteration I	5
The number of CPU cycles to process one sample C	7×10^5
The effective switch capacitance constant ζ	1×10^{-28}
The maximum CPU frequency of device n $f_{n,\max}$	[0.5, 4.0] GHz
The maximum transmission power of device n $p_{n,\max}$	[0.001, 1] W
The total bandwidth of the system B	2 MHz
Channel coefficient of device n during the t -h round $G_{t,n}$	[10^{-7} , 10^{-6}] dB
The noise power spectral density N_0	-174 dbm/MHz
Weight size of the local model in bits	53.21 Mbit
Time-related QoS T_{QoS}	15 s
Weight factor σ	0.8

consumption can be known in advance, so that the ϵ -greedy algorithm can be executed [7].

- FL-DLT3 [25]: FL-DLT3 enables a twin-delayed deep deterministic policy gradient (TD3) framework to optimize accuracy and energy balance in a continuous domain. Compared to transmission power allocation for FL efficiency optimization [25], we further expand it to manage CPU frequency with equal, fixed bandwidth allocation.
- SAC (Soft Actor-Critic) [10], [46]: SAC is an off-policy actor-critic deep RL algorithm based on the maximum entropy framework [46]. [10] applies SAC to wireless FL scenarios for resource allocation. We implement SAC on top of a transformer decoder-only neural network.
- T-ELLM PPO: It uses PPO [41] rather than GRPO [38] to fine-tune T-ELLM. Compared to GRPO, PPO employs an additional critic network to optimize the NN.

B. Performance Comparison

We first show the overall performance comparison, in terms of CWEPM in (8a), learning accuracy, and energy consumption per round, and Fig. 5 presents the corresponding results. Notably, as shown in Fig. 6, significantly heterogeneous label distributions and dataset sizes exist for evaluation. As shown in Fig. 5(a), the proposed T-ELLM GRPO and




Fig. 5. Performance comparison of different algorithms for wireless FL.




Fig. 6. Label and dataset size distributions for non-IID data on different devices.

T-ELLM PPO achieve the highest CWPEM values, suggesting that these methods require less energy to attain target performance levels. Fig. 5(b) illustrates the convergence of wireless FL under different algorithms. The proposed T-ELLM GRPO and T-ELLM PPO exhibit a remarkable convergence rate towards the desired accuracy. Additionally, Fig. 5(c) presents the energy consumption per round of wireless FL, where the solid curves therein represent the mean testing accuracy across 10 experimental trials and the shaded part represents the 95% confidence interval calculated from these experiments (the other figures are likewise). The proposed T-ELLM GRPO and T-ELLM PPO consume less energy per round, demonstrating their efficiency in resource utilization during the FL process. Furthermore, the narrower confidence interval observed for T-ELLM GRPO and T-ELLM PPO, relative to the FL-DLT3 and Fedavg Tool baseline, indicates stable and consistent decision-making, while maintaining a level of stability comparable to the heuristic methods.

To implement the proposed T-ELLM, 600 trajectories of $((s_t^m, \Xi(S_{t-1}), S_t), (s_{t+1}^m, \Xi(S_t)))$ are collected to train the GPT model as the statistics part of the virtual environment. The convergence of the model-based environment is shown in Fig. 7(a). It can be observed from Fig. 7(a) that the MSE loss for both state s_t^m and accuracy $\Xi(S_t)$ decreases steadily.

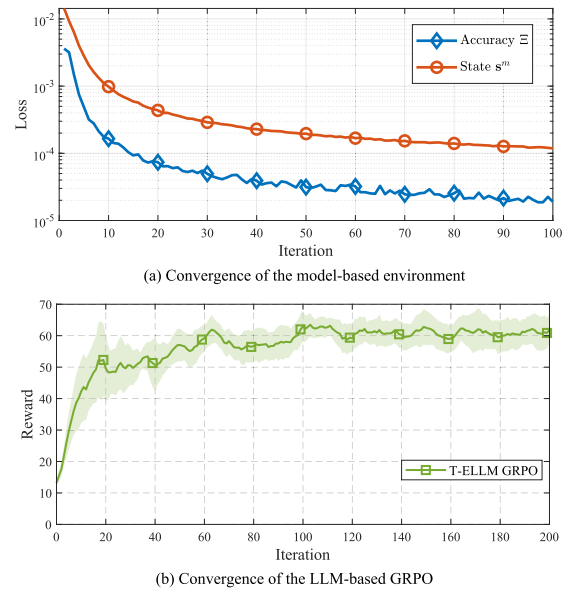


Fig. 7. Convergence of the model-based virtual environment and GRPO-based T-ELLM.




Fig. 8. Comparison of changing environment with respect to $f_{n,max}$ and $P_{n,max}$.

While the accuracy metric consists of a single value, thus being relatively easier to predict, the state encompasses more

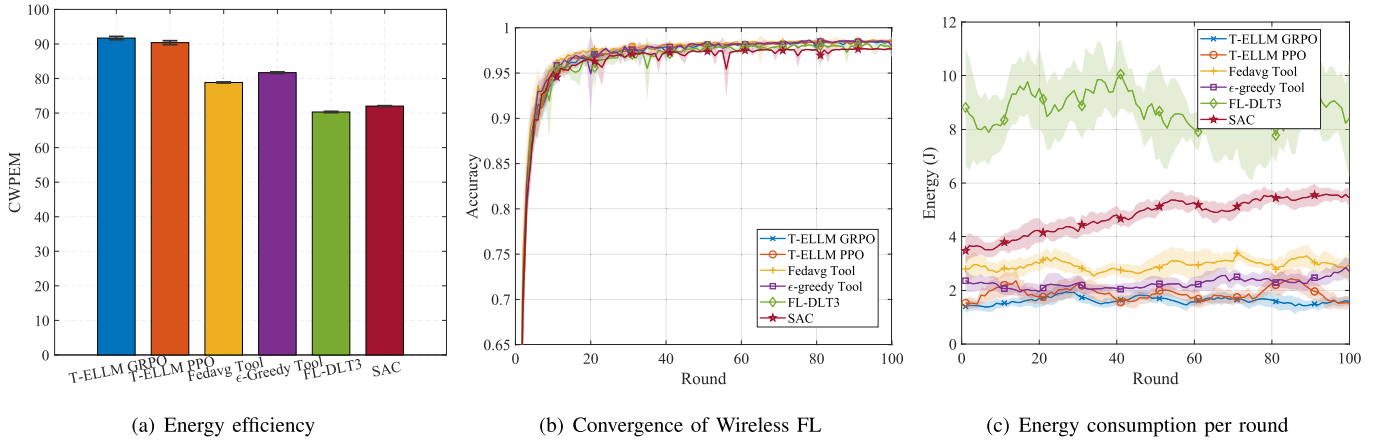


Fig. 9. Performance comparison of different algorithms for wireless FL with Scenario 2.




Fig. 10. Performance comparison of different algorithms for wireless FL with $T_{QoS} = 20$ s.

complex information. Nevertheless, the state also converges to a satisfactory result. Based on the trained environmental model, the evolution of T-ELLM can be carried out offline. Fig. 7(b) presents the reward progression during GRPO fine-tuning. It can be seen that under the optimization of GRPO, the reward value gradually stabilizes at a higher level, which indicates the satisfactory convergence of T-ELLM.

C. Generalization Ability of T-ELLM

1) *Adaptation to Resource Changes:* To verify the generalization ability of the proposed T-ELLM in response to the changes in the environment, we first change the computing and communication capabilities of each device in the environment. Fig. 8 highlights the differences in terms of the maximum CPU frequency and transmission power of devices. Compared to the environment (i.e., Scenario 1) for evaluation in Section VI-B, the modified environment is termed Scenario 2. Notably, all algorithms are directly transferred from the training outcome in Scenario 1 and have no prior interaction with Scenario 2, ensuring an unbiased assessment of adaptability. The corresponding performance in Scenario 2 is shown in Fig. 9. As demonstrated in Fig. 9(a), the CWPEM of T-ELLM GRPO and T-ELLM PPO outperforms that of other baseline algorithms, indicating superior performance with lower energy

consumption. Furthermore, Fig. 9(b) confirms that the learning accuracy of T-ELLM remains robust, without sacrificing the learning efficiency. Furthermore, Fig. 9(c) shows that the energy consumption per round yielded by T-ELLM is also lower than that of other algorithms, demonstrating its ability to maintain high accuracy while optimizing energy efficiency. These results highlight T-ELLM's adaptability to varying environmental conditions. Additionally, the FedAvg tool and ϵ -Greedy tool are not learning-based algorithms, naturally, their performance is less affected by environmental changes. In comparison, as shown in Fig. 9(c), due to its training only in the default Scenario 1, the energy consumption of FL-DLT3 and SAC is significantly higher than that of other algorithms. This further underscores the advantages of the proposed T-ELLM in terms of environmental adaptability and energy efficiency.

2) *Adaptation to Changed Task Requirements:* In this part, we further demonstrate the ability of T-ELLM to cope with different QoS requirements. Fig. 10 provides the results after changing T_{QoS} in the default Scenario 1 from 15 seconds to 20 seconds. The results show that the proposed T-ELLM GRPO and T-ELLM PPO achieve the highest CWPEM values in Fig. 10(a), indicating superior energy efficiency in wireless




Fig. 11. Performance of the proposed T-ELLM GRPO with different QoS requirements and data heterogeneity.

FL. Fig. 10(b) illustrates the convergence of wireless FL under different algorithms. The proposed T-ELLM GRPO and T-ELLM PPO exhibit a significant convergence rate, achieving the desired accuracy. Fig. 10(c) shows that the proposed T-ELLM GRPO and T-ELLM PPO consume less energy per round, while it can be seen that the FedAvg Tool also maintains a low energy consumption per round. The reason lies in that the resource tool is used therein for resource allocation, contributing to saving overall energy consumption required for the anticipated QoS. On the other hand, although the ϵ -Greedy tool algorithm also gets help from the tool, the extensive reliance on the artificial setting of greedy strategy parameters undermines the potential benefit. Additionally, as shown in Fig. 10(c), FL-DLT3 and SAC demonstrate notably higher energy consumption, as they depend solely on their NNs for resource allocation and lack adaptability to varying QoS demands. In comparison, the proposed T-ELLM framework dynamically adjusts to QoS requirements, enabling more efficient decision-making.

We also evaluate the performance of the proposed T-ELLM with different QoS requirements and heterogeneity. The performance of the proposed T-ELLM is shown in Fig. 11. It can be seen in Fig. 11(a) that the proposed T-ELLM maintains comparable accuracy with different QoS requirements. Furthermore, the energy consumption decreases with the increase of QoS requirements. Fig. 11(b) shows the performance of the proposed T-ELLM with different levels of data heterogeneity, and similar observations can be attained. These results show that the proposed T-ELLM can adapt to environmental changes by yielding appealing accuracy.

D. Performance of Tellm in Different Model Sizes

In order to validate the effect of LLM size and language training approach on the T-ELLM, we evaluate the performance of three versions of GPT-2, including GPT-2 Small, GPT-2 Medium, and GPT-2 Large, corresponding to parameter

scales of 0.1B, 0.3B, and 0.7B, respectively. Additionally, we compare the performance of GPT-2 models against a more lightweight, task-specific baseline we call GPT-Linear, which uses a Transformer decoder-only architecture (like GPT-2). The result is shown in Fig. 12. It can be seen that the performance for the GPT-2 series is all very similar, with no significant difference. Compared with the GPT-2 series, the CWPEM and convergence results of GPT-Linear are slightly lower, while the confidence interval of GPT-Linear is narrower than that of the GPT-2 series. These results validate the advantage of adopting a linguistic model, and a computationally efficient GPT-2 Small model already yields satisfactory results.

E. Adaptability of Tellm to Different FL Algorithms

We investigate the adaptability and robustness of our T-ELLM framework for FL with gradient quantization [10] and Differential Privacy (DP) [4].

1) *Adaptation to Gradient Quantization*: We present the performance of T-ELLM with varying levels of gradient quantization in Fig. 13. As shown in Fig. 13(a), consistent with our intuition, the energy efficiency of T-ELLM decreases as the number of quantization bits decreases, which is expected since lower bit quantization introduces more noise, potentially degrading model accuracy. However, even with 2-bit quantization, T-ELLM maintains a reasonable balance between energy consumption and accuracy. Fig. 13(b) shows that T-ELLM converges to a satisfactory accuracy level across all quantization levels, although higher bit quantization leads to faster convergence and higher final accuracy. Fig. 13(c) indicates that energy consumption per round increases with the number of quantization bits, as higher precision requires more data transmission. Overall, T-ELLM demonstrates robust performance even under aggressive gradient quantization.

2) *Adaptation to Differential Privacy*: We implement (ϵ, δ) -DP in the wireless FL process, where δ accounts for the probability that plain ϵ -DP is broken, and ϵ controls the level of privacy, where a smaller ϵ signifies stronger privacy (and more noise). We evaluate the performance with various privacy budgets. Fig. 14(a) shows that stronger privacy (a lower ϵ) leads to lower final model accuracy, as the increased noise makes learning more difficult. Fig. 14(b) demonstrates that while all settings eventually reach a stable accuracy, stronger privacy constraints slow down the learning process. Fig. 14(c) illustrates that energy consumption is not significantly affected by the DP level. Overall, T-ELLM maintains robust performance when differential privacy is applied.

F. Scalability for Large-Scale Scenario

To evaluate the scalability of T-ELLM, we simulate a larger wireless FL environment consisting of 100 devices, where we distribute the CIFAR-10 dataset across these devices and configure the task to select 10 participating devices in each training round. Meanwhile, we compare the performance of three algorithms:

- T-ELLM Retrain: The T-ELLM agent is retrained using the virtual environment for the 100-device scenario.




Fig. 12. Performance comparison of large models with different scales for wireless FL.




Fig. 13. Performance of the T-ELLM algorithm for wireless FL using gradient quantization.




Fig. 14. Performance of the T-ELLM algorithm for wireless FL using differential privacy.

- **T-ELLM Original:** The original T-ELLM agent, trained for the 20-device scenario, was directly applied without any retraining to test its zero-shot transferability.
- **SAC-Transformer:** A baseline using a Transformer architecture with a linear head, trained from scratch on the 100-device scenario using the SAC algorithm.

The results of these experiments are presented in Fig. 15. As shown in Fig. 15(a), the T-ELLM Retrain model achieves the highest energy efficiency and CWPEM value. Meanwhile,

the T-ELLM Original model also performs well, demonstrating strong transferability. The SAC-Transformer baseline is comparatively less efficient. Fig. 15(b) shows that both T-ELLM Retrain and T-ELLM Original achieve higher convergence accuracy than SAC-Transformer. In other words, the T-ELLM Original model can directly make effective device selections in a much larger environment. Finally, Fig. 15(c) illustrates that the T-ELLM Retrain is the most energy-efficient. On the contrary, the SAC-Transformer consumes drastically more




Fig. 15. Performance comparison of learning-based algorithms migrated to the scenario with 100 devices.

energy. These results indicate that the T-ELLM framework provides a robust and scalable solution for large-scale FL resource management.

VII. CONCLUSION

We have presented T-ELLM, a novel tool-aided evolutionary LLM framework for efficient device selection and resource allocation in wireless FL. Notably, T-ELLM has been developed on top of a mathematics-driven decoupling of LLM-based device selection and tool-based resource allocation. Meanwhile, the combination of the linguistic reasoning capabilities in LLMs and mathematical optimization tools contributes to boosting the generalization capability of decision-making in environmental changes. In addition, T-ELLM takes advantage of a model-based virtual environment to support GRPO-based fine-tuning at minimal communication cost during interactions. Our theoretical analysis has proved the bounded discrepancy between virtual and real environments, which ensures the advantage function learned in the virtual environment maintains a provably small deviation from real-world conditions. Extensive experimental results have demonstrated that T-ELLM can further improve the energy efficiency and exhibit robust adaptability to environmental changes. The T-ELLM framework can be extended to more complicated wireless resource management tasks, such as task offloading [47], [48]. Other potential directions include investigation of its scalability, enhancing its safety against attacks, and validating its performance on real-world hardware. Also, its algorithmic complexity can be refined for practical deployment.

REFERENCES

- [1] Q. Cui, "Overview of AI and communication for 6G network: Fundamentals, challenges, and future research opportunities," *Sci. China Inf. Sci.*, vol. 68, no. 7, pp. 1–61, Jan. 2025, Art. no. 171301.
- [2] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2016, pp. 1273–1282.
- [3] H. Guan, P.-T. Yap, A. Bozoki, and M. Liu, "Federated learning for medical image analysis: A survey," *Pattern Recognit.*, vol. 151, Jul. 2024, Art. no. 110424.
- [4] M. Ye, X. Fang, B. Du, P. C. Yuen, and D. Tao, "Heterogeneous federated learning: State-of-the-art and research challenges," *ACM Comput. Surveys*, vol. 56, no. 3, pp. 1–44, Oct. 2023.
- [5] H. T. Nguyen, V. Schwag, S. Hosseinalipour, C. G. Brinton, M. Chiang, and H. Vincent Poor, "Fast-convergent federated learning," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 1, pp. 201–218, Jan. 2021.
- [6] X. Yi, R. Li, C. Peng, F. Wang, J. Wu, and Z. Zhao, "RHFedMTL: Resource-aware hierarchical federated multitask learning," *IEEE Internet Things J.*, vol. 11, no. 14, pp. 25227–25238, Jul. 2024.
- [7] F. Lai, X. Zhu, H. V. Madhyastha, and M. Chowdhury, "Oort: Efficient federated learning via guided participant selection," in *Proc. USENIX Symp. Oper. Syst. Des. Implement. (OSDI)*, Jul. 2021, pp. 19–35.
- [8] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in *Proc. ICC - IEEE Int. Conf. Commun. (ICC)*, Shanghai, China, May 2019, pp. 1–7.
- [9] S. Q. Zhang, J. Lin, and Q. Zhang, "A multi-agent reinforcement learning approach for efficient client selection in federated learning," in *Proc. AAAI Conf. Artif. Intell.*, Jun. 2022, vol. 36, no. 8, pp. 9091–9099.
- [10] X. Hou, J. Wang, C. Jiang, Z. Meng, J. Chen, and Y. Ren, "Efficient federated learning for metaverse via dynamic user selection, gradient quantization and resource allocation," *IEEE J. Sel. Areas Commun.*, vol. 42, no. 4, pp. 850–866, Apr. 2024.
- [11] Y. G. Kim and C.-J. Wu, "AutoFL: Enabling heterogeneity-aware energy efficient federated learning," in *Proc. Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, Oct. 2021, pp. 183–198.
- [12] H. Wang, Z. Kaplan, D. Niu, and B. Li, "Optimizing federated learning on non-IID data with reinforcement learning," in *Proc. IEEE Conf. Comput. Commun.*, Jul. 2020, pp. 1698–1707.
- [13] Y. Zhan, P. Li, and S. Guo, "Experience-driven computational resource allocation of federated learning by deep reinforcement learning," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. (IPDPS)*, New Orleans, LA, USA, May 2020, pp. 234–243.
- [14] R. Figueiredo Prudencio, M. R. O. A. Maximo, and E. L. Colombari, "A survey on offline reinforcement learning: Taxonomy, review, and open problems," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 8, pp. 10237–10257, Aug. 2024.
- [15] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Comput. Surv.*, vol. 50, no. 2, pp. 1–35, Apr. 2017.
- [16] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," in *Handbook of Reinforcement Learning and Control*. Cham, Switzerland: Springer, Jun. 2021, pp. 321–384.
- [17] J. Li, T. Chen, and S. Teng, "A comprehensive survey on client selection strategies in federated learning," *Comput. Netw.*, vol. 251, Sep. 2024, Art. no. 110663.
- [18] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A joint learning and communications framework for federated learning over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 269–283, Jan. 2021.
- [19] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Babaei, "Energy efficient federated learning over wireless communication networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1935–1949, Mar. 2021.


- [20] J. Yao and N. Ansari, "Enhancing federated learning in fog-aided IoT by CPU frequency and wireless power control," *IEEE Internet Things J.*, vol. 8, no. 5, pp. 3438–3445, Mar. 2021.
- [21] Y. Jee Cho, J. Wang, and G. Joshi, "Towards understanding biased client selection in federated learning," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, Mar. 2022, pp. 10351–10375.
- [22] W. Shi, S. Zhou, Z. Niu, M. Jiang, and L. Geng, "Joint device scheduling and resource allocation for latency constrained wireless federated learning," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 453–467, Jan. 2021.
- [23] T. Zhang, K.-Y. Lam, J. Zhao, and J. Feng, "Joint device scheduling and bandwidth allocation for federated learning over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 24, no. 1, pp. 3–18, Jan. 2025.
- [24] W. Mao, X. Lu, Y. Jiang, and H. Zheng, "Joint client selection and bandwidth allocation of wireless federated learning by deep reinforcement learning," *IEEE Trans. Services Comput.*, vol. 17, no. 1, pp. 336–348, Jan. 2024.
- [25] J. Zheng, K. Li, N. Mhaisen, W. Ni, E. Tovar, and M. Guizani, "Exploring deep-reinforcement-learning-assisted federated learning for online resource allocation in privacy-preserving EdgeIoT," *IEEE Internet Things J.*, vol. 9, no. 21, pp. 21099–21110, Nov. 2022.
- [26] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. (Jun. 2018). *Improving Language Understanding By Generative Pre-Training*. OpenAI Blog. [Online]. Available: <https://cdn.openai.com/research-covers/language-unsupervised/languageunderstandingpaper.pdf>
- [27] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. (Aug. 2019). *Language Models Are Unsupervised Multitask Learners*. [Online]. Available: <https://cdn.openai.com/better-language-models/language-models-are-unsupervised-multitask-learners.pdf>
- [28] A. Grattafiori et al., "The llama 3 herd of models," 2024, *arXiv:2407.21783*.
- [29] A. Liu et al., "DeepSeek-V3 technical report," 2024, *arXiv:2412.19437*.
- [30] S. Reed et al., "A generalist agent," *Transact. Mach. Learn. Res.*, pp. 1–42, Nov. 2022. [Online]. Available: <https://openreview.net/forum?id=likK0kHjvj>
- [31] Z. Wang, Y. Zhou, Y. Shi, and K. B. Letaief, "Federated fine-tuning for pre-trained foundation models over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 24, no. 4, pp. 3450–3464, Apr. 2025.
- [32] J. Zheng, H. Zhang, L. Wang, W. Qiu, H. Zheng, and Z. Zheng, "Safely learning with private data: A federated learning framework for large language model," in *Proc. Conf. Empir. Methods Nat. Lang. Process. (EMNLP)*, Nov. 2024, pp. 5293–5306.
- [33] C. E. Mower et al., "ROS-LLM: A ROS framework for embodied AI with task feedback and structured reasoning," 2024, *arXiv:2406.19741*.
- [34] M. T. R. Laskar et al., "A systematic survey and critical review on evaluating large language models: Challenges, limitations, and recommendations," in *Proc. Conf. Empir. Methods Nat. Lang. Process. (EMNLP)*, Nov. 2024, pp. 13785–13816.
- [35] J. Ahn, R. Verma, R. Lou, D. Liu, R. Zhang, and W. Yin, "Large language models for mathematical reasoning: Progresses and challenges," in *Proc. Eur. Chapter Assoc. Comput. Linguist. (EACL)*, Mar. 2024, pp. 225–237.
- [36] D. Bandyopadhyay, S. Bhattacharjee, and A. Ekbal, "Thinking machines: A survey of LLM based reasoning strategies," 2025, *ArXiv:2503.10814*.
- [37] M. Ahn et al., "Do as I can, not as I say: Grounding language in robotic affordances," in *Proc. 6th Conf. Robot Learn.*, vol. 205, Dec. 2023, pp. 287–318.
- [38] Z. Shao et al., "DeepSeekMath: Pushing the limits of mathematical reasoning in open language models," 2024, *arXiv:2402.03300*.
- [39] N. H. Tran, W. Bao, A. Zomaya, M. N. H. Nguyen, and C. S. Hong, "Federated learning over wireless networks: Optimization model design and analysis," in *Proc. IEEE Conf. Comput. Commun.*, Apr. 2019, pp. 1387–1395.
- [40] A. Albaseer, M. Abdallah, A. Al-Fuqaha, and A. Erbad, "Data-driven participant selection and bandwidth allocation for heterogeneous federated edge learning," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 53, no. 9, pp. 5848–5860, Sep. 2023.
- [41] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [42] L. Devroye, A. Mehrabian, and T. Reddad, "The total variation distance between high-dimensional Gaussians with the same mean," 2018, *arXiv:1810.08693*.
- [43] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. Providence, RI, USA: American Mathematical Society, 2009.
- [44] M. Kearns and S. Singh, "Near-optimal reinforcement learning in polynomial time," *Mach. Learn.*, vol. 49, nos. 2–3, pp. 209–232, Nov. 2002.
- [45] L. Chen et al., "Decision transformer: Reinforcement learning via sequence modeling," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, Dec. 2021, pp. 15084–15097.
- [46] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1861–1870.
- [47] M. Asif, F. M. Aldosari, D. Y. Badawood, and A. Salam, "CAMP-GNN: A constraint-aware message-passing model for optimal resource allocation in software projects," *IEEE Access*, vol. 13, pp. 87965–87977, 2025.
- [48] M. Asif et al., "Advanced zero-shot learning (AZSL) framework for secure model generalization in federated learning," *IEEE Access*, vol. 12, pp. 184393–184407, 2024.



Chongyang Tan received the M.S. degree from Jinan University, Guangzhou, China, in 2024. He is currently pursuing the Ph.D. degree with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China. His research interests include wireless communication, large language models, distributed learning, and reinforcement learning.



Ruoqi Wen received the B.Sc. degree in mathematics and applied mathematics from Hefei University of Technology, Hefei, China, in June 2021. She is currently pursuing the Ph.D. degree with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China. Her research interests include multi-agent reinforcement learning, regret bounds, endogenous and exogenous uncertainty analysis, wireless networks, and autonomous vehicle control.



Rongpeng Li (Senior Member, IEEE) received the B.E. degree from Xidian University, Xi'an, China, in June 2010, and the Ph.D. degree from Zhejiang University, Hangzhou, China, in June 2015. From August 2015 to September 2016, he was a Research Engineer with the Wireless Communication Laboratory, Huawei Technologies Company Ltd., Shanghai, China. He was a Visiting Scholar with the Department of Computer Science and Technology, University of Cambridge, Cambridge, U.K., from February 2020 to August 2020. He is currently an Associate Professor with the College of Information Science and Electronic Engineering, Zhejiang University. His current research interests include networked intelligence for comprehensive efficiency (NICE).



Zhifeng Zhao (Member, IEEE) received the B.E. degree in computer science, the M.E. degree in communication and information systems, and the Ph.D. degree in communication and information systems from the PLA University of Science and Technology, Nanjing, China, in 1996, 1999, and 2002, respectively. From 2002 to 2004, he was a Post-Doctoral Researcher with Zhejiang University, Hangzhou, China, where his researches were focused on multimedia next-generation networks (NGNs) and softswitch technology for energy efficiency.

Currently, he is with Zhejiang Laboratory, Hangzhou, as the Chief Engineering Officer. His research interests include software-defined networks (SDNs), wireless networks in 6G, computing networks, and collective intelligence. He is the Symposium Co-Chair of ChinaCom 2009 and 2010. He was the Technical Program Committee (TPC) Co-Chair of the Tenth IEEE International Symposium on Communication and Information Technology (ISCIT 2010).



Ekram Hossain (Fellow, IEEE) is currently a Professor and the Associate Head (Graduate Studies) of the Department of Electrical and Computer Engineering, University of Manitoba, Canada. He is a member (Class of 2016) of the College of the Royal Society of Canada. He is a fellow of Canadian Academy of Engineering and the Engineering Institute of Canada. He has won several research awards, including the 2017 IEEE Communications Society Best Survey Paper Award and the 2011 IEEE Communications Society Fred Ellersick Prize Paper Award.

He was listed as a Clarivate Analytics Highly Cited Researcher in computer science from 2017 to 2025. Previously, he served as the Editor-in-Chief for IEEE Press from 2018 to 2021 and IEEE COMMUNICATIONS SURVEYS AND TUTORIALS from 2012 to 2016. He was a Distinguished Lecturer of the IEEE Communications Society and the IEEE Vehicular Technology Society. He served as the Director of Magazines and the Director of Online Content for the IEEE Communications Society from 2020 to 2021 and from 2022 to 2023, respectively. He is the IEEE Canada President-Elect from 2026 to 2027.



Honggang Zhang (Fellow, IEEE) was a Professor with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China. He was an Honorary Visiting Professor with the University of York, York, U.K., and an International Chair Professor of Excellence with the Université Européenne de Bretagne, Supélec, France. He is currently a Professor with the School of Computer Science and Engineering, Macau University of Science and Technology, Macau, China.

His research interests include cognitive radio networks, semantic communications, green communications, machine learning, artificial intelligence, intelligent computing, and the internet of intelligence.

Dr. Zhang was a co-recipient of the 2021 IEEE Communications Society Outstanding Paper Award and the 2021 IEEE Internet of Things Journal (IOT-J) Best Paper Award. He served as the Chair for the Technical Committee on Cognitive Networks of the IEEE Communications Society from 2011 to 2012. He was the Founding Chief Managing Editor of *Intelligent Computing* (Science Partner Journal). He was the leading Guest Editor of the Special Issues on Green Communications for *IEEE Communications Magazine*. He served as a Series Editor for *IEEE Communications Magazine* (Green Communications and Computing Networks Series) from 2015 to 2018. He is the Associate Editor-in-Chief of *China Communications*.