

Incentivizing Efficient Zeroth-Order Federated Fine-Tuning of Generative AI in Satellite-UAV Cooperative Networks

Haoqing Jiang¹, Qichao Xu¹, Zhou Su², Marie Siew³, Yihao Qi¹, Zehui Xiong⁴

¹School of Mechatronic Engineering and Automation, Shanghai University, Shanghai, China

²School of Cyber Science and Engineering, Xi'an Jiaotong University, Xi'an, China

³Information Systems Technology and Design, Singapore University of Technology and Design, Singapore

⁴School of Electronics, Electrical Engineering and Computer Science (EEECS), Queen's University Belfast, U.K

Abstract—Generative artificial intelligence (GenAI) provides new opportunities for unmanned aerial vehicles (UAVs) to perform intelligent missions such as disaster response in remote areas. Since the on-board observations are often incomplete in such harsh environments, GenAI synthesizes realistic sensory data to augment the local datasets, which enhances the perception and decision-making of UAVs. Satellite-based federated fine-tuning further enables collaborative GenAI training without sharing raw data or relying on terrestrial infrastructures. However, the limited satellite bandwidth and onboard computation resources, together with self-interested UAVs holding private computation capabilities, pose significant challenges to efficient and incentivized GenAI fine-tuning. In this paper, we propose a game theoretic zeroth order (ZO) federated fine-tuning framework for GenAI in satellite-UAV networks. Specifically, the ZO optimization is adopted, which replaces the high-dimensional gradient uploading with a few scalar projections. Then, we formulate the query-reward contract design problem between the satellite and the UAVs. To deal with the private computation information of UAVs, we design an incentive compatible contract mechanism, which motivates each UAV to truthfully select the optimal contract item designed for its own type. Simulation results demonstrate that the proposed scheme can effectively improve the efficiency of satellite assisted GenAI fine-tuning.

Index Terms—Generative artificial intelligence (GenAI), federated fine-tuning, satellite networks, contract theory.

I. INTRODUCTION

WITH the development of generative artificial intelligence (GenAI), unmanned aerial vehicles (UAVs) have been applied to intelligent missions such as environment exploration and disaster response [1]. Benefiting from the empowerment of GenAI, UAVs can synthesize realistic visual scenes of the unexplored areas and reason about the environment from limited observations, which supports intelligent decisions such as identifying the affected regions and planning where to dispatch emergency services. In particular, the UAVs can be flexibly deployed to hard-to-reach areas where ground access and infrastructures are unavailable [2]. However, such GenAI capabilities rely on continual model training, which is hard to sustain in these areas.

Recently, satellite communication networks provide a feasible solution to the above dilemma, where UAVs can rely on

the wide-coverage satellite links to accomplish critical mission services [3]. To collaboratively fine-tuning the global GenAI model coordinated by the satellite without sharing the raw data, federated fine-tuning offers a promising paradigm [4]. Moreover, the seed-based zeroth-order (ZO) optimization has emerged as a communication-efficient training approach, where each model owner only performs forward evaluations of the loss function and uploads a few scalar projections instead of the high-dimensional gradients [5].

However, directly applying the ZO federated fine-tuning to the considered scenario still faces several challenges. First, performing the ZO queries consumes the limited onboard computation resources and energy of UAVs, which restricts the training contribution that each UAV can afford [6]. Second, due to the orbital motion, each UAV is only intermittently visible to the satellite within limited contact windows, which further constrains the training that can be accomplished per round [7]. Third, the UAVs are self-interested individuals and hold private computation capabilities that are unknown to the satellite, so that an effective incentive mechanism is needed to motivate their truthful participation under such information asymmetry [8]. However, how to jointly handle the resource constraints and motivate the truthful participation of UAVs remains an open issue.

In this paper, we propose a game theoretic ZO federated fine-tuning framework for GenAI in satellite-UAV networks. The main contributions of this paper are summarized as follows. Firstly, in the proposed framework, a low Earth orbit (LEO) satellite serves as the parameter server and UAVs act as distributed clients, which upload only scalar projections instead of high dimensional gradients via the seed-based ZO optimization to reduce the communication overhead. Secondly, we characterize the training contribution of each UAV by the number of ZO queries, and model the resource constraints, including the channel condition, the computation and communication costs, and the satellite visibility window. Thirdly, we formulate the interactions between the satellite and heterogeneous UAVs as a query-reward contract theory under information asymmetry, and design an incentive-compatible mechanism to motivate each UAV to truthfully select the contract item designed for its type. Finally, simulation results demonstrate that the proposed scheme improves the fine-tuning performance of the GenAI model and achieves higher utilities.

This paper is supported by the Key Industry R&D Program of Shenzhen (Grant No. ZDCY20250901112700001), NSFC No. 62371280 and A*STAR under its IAF-ICP programme (H25-MCP3438). Corresponding Author: Qichao Xu (xqc690926910@shu.edu.cn) and Marie Siew (marie_siew@sutd.edu.sg)




Fig. 1. The satellite-UAV ZO federated GenAI fine-tuning model

II. SYSTEM MODEL

A. Network Model

As shown in Fig. 1, we consider a satellite-UAV federated diffusion learning system, which consists of one LEO satellite serving as the parameter server and a group of UAVs serving as clients. The LEO satellite, denoted by o , coordinates the UAVs to collaboratively train a generative diffusion model over their local data. Since the diffusion model is large and full-model training is infeasible on resource-constrained UAVs, we adopt a parameter-efficient fine-tuning scheme: the pre-trained backbone $\mathbf{w}_0 \in \mathbb{R}^D$ is frozen, and only a set of lightweight adapter parameters $\phi \in \mathbb{R}^{D_a}$ is trained, where $D_a \ll D$ is the number of trainable parameters [4]. Specifically, for each modulated layer ℓ of the backbone with C_ℓ output channels, FiLM applies the per channel affine transformation $h_{\ell,c} \leftarrow \gamma_{\ell,c} h_{\ell,c} + \beta_{\ell,c}$, where only the scale and shift parameters $\{\gamma_{\ell,c}, \beta_{\ell,c}\} \in \mathbb{R}^{C_\ell}$ are trained while the backbone weights stay fixed, so that $D_a = 2 \sum_\ell C_\ell$. Let $\mathcal{I} = \{1, \dots, i, \dots, I\}$ denote the set of UAVs, where each UAV i is mounted with sensors and dispatched to a task area to execute a specific mission, and collects a local dataset $\mathcal{D}_i$ during the mission. The adapter parameters are trained by minimizing the denoising score-matching loss over all UAVs [9],

$$\min_{\phi \in \mathbb{R}^{D_a}} F(\phi) = \sum_{i \in \mathcal{I}} \frac{|\mathcal{D}_i|}{\sum_k |\mathcal{D}_k|} F_i(\phi), \quad (1)$$

where the local denoising loss of UAV i is

$$F_i(\phi) = \mathbb{E}_{\tau, \mathbf{x}_0 \sim \mathcal{D}_i, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \epsilon - \epsilon_{\mathbf{w}_0, \phi}(\mathbf{x}_\tau, \tau) \right\|^2 \right]. \quad (2)$$

Here, τ is a diffusion time step (noise scale), $\mathbf{x}_\tau = \sqrt{\bar{\alpha}_\tau} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_\tau} \epsilon$ is the noised sample with the noise schedule $\{\bar{\alpha}_\tau\}$, ϵ is the added Gaussian noise, and $\epsilon_{\mathbf{w}_0, \phi}(\cdot, \cdot)$ is the noise-prediction network with the frozen backbone $\mathbf{w}_0$ and the trainable adapter parameters ϕ . In each loss evaluation, a diffusion time step τ is sampled and the denoising residual at that step is computed.

B. Zeroth-Order Diffusion Model

To enable on device training of the adapter parameters under the limited onboard resources of UAVs, each UAV employs a seed-based ZO optimizer that avoids backpropagation. The ZO optimizer estimates the gradient with respect to ϕ using only

forward evaluations of the denoising loss $F_i(\phi)$ in (2), where each evaluation is termed a query.

Specifically, in each training round, the satellite broadcasts the adapter parameters ϕ together with a random seed ξ to the UAVs, while the frozen backbone $\mathbf{w}_0$ was pre-loaded on the UAVs. Upon reception, UAV i regenerates a set of perturbation vectors $\{\mathbf{u}_{i,m}\}_{m=1}^{b_i}$ from the seed ξ , where each $\mathbf{u}_{i,m} \in \mathbb{R}^{D_a}$ is the m -th perturbation direction sampled from a standard Gaussian distribution in the low dimensional adapter space and b_i is the number of perturbation directions. For each direction $\mathbf{u}_{i,m}$, UAV i computes a two-point gradient estimate of F_i with respect to ϕ as

$$\hat{g}_{i,m} = \frac{F_i(\phi + \mu \mathbf{u}_{i,m}) - F_i(\phi - \mu \mathbf{u}_{i,m})}{2\mu} \mathbf{u}_{i,m}, \quad (3)$$

where $\mu > 0$ is the smoothing parameter and only the adapter parameters are perturbed while the backbone $\mathbf{w}_0$ stays fixed. Here, computing each two-point estimate in (3) requires evaluating the loss function twice, and hence the number of queries of UAV i is defined as $K_i = 2b_i$. The ZO gradient estimate of UAV i is then averaged over the b_i directions as $\hat{\mathbf{g}}_i = \frac{1}{b_i} \sum_{m=1}^{b_i} \hat{g}_{i,m}$. After local computation, UAV i uploads only the b_i scalar projections $\{(F_i(\phi + \mu \mathbf{u}_{i,m}) - F_i(\phi - \mu \mathbf{u}_{i,m})) / (2\mu)\}_{m=1}^{b_i}$ to the satellite, and the satellite reconstructs $\hat{\mathbf{g}}_i$ using the shared seed ξ . Finally, the satellite aggregates the gradient estimates from all UAVs to update the adapter parameters as $\phi \leftarrow \phi - \gamma \sum_i \frac{|\mathcal{D}_i|}{\sum_k |\mathcal{D}_k|} \hat{\mathbf{g}}_i$, where γ is the learning rate.

The number of queries K_i characterizes the local training effort of UAV i , since a larger K_i reduces the variance of the ZO gradient estimate and thus accelerates the convergence, with a diminishing return of the order $\mathcal{O}(1/\sqrt{K_i})$ on the convergence error [5]. Besides, the uplink payload of UAV i is only $b_i = K_i/2$ scalars, which is independent of the model dimension.

C. Mobility Model

To characterize the movement of the UAVs and the satellite, the total time horizon is evenly divided into T time slots with length Δt . The instant location of UAV i at time slot t is denoted as $\mathbf{r}_i(t) = [x_i(t), y_i(t), z_i(t)]$, where $[x_i(t), y_i(t)]$ is the horizontal coordinate and $z_i(t)$ is the hovering altitude. The trajectory of each UAV is predetermined by the mission. Meanwhile, the LEO satellite moves along a deterministic orbit with an orbital period T_{orb} , and its instant position at time slot t is denoted as $\mathbf{r}_{\text{sat}}(t)$, which can be accurately predicted from the orbital parameters. As such, the distance between UAV i and the satellite o at time slot t is

$$d_i(t) = \|\mathbf{r}_{\text{sat}}(t) - \mathbf{r}_i(t)\|, \quad (4)$$

and the elevation angle of the satellite o seen from UAV i is

$$\varphi_i(t) = \arcsin\left(\frac{z_{\text{sat}}(t) - z_i(t)}{d_i(t)}\right), \quad (5)$$

where $z_{\text{sat}}(t)$ is the altitude of the satellite o , and a larger elevation angle $\varphi_i(t)$ implies a shorter and less obstructed link between UAV i and the satellite o .

D. Channel Model

The UAV-to-satellite uplink adopts the line-of-sight (LoS) dominant channel model. The path loss between UAV i and the satellite o at time slot t is

$$\Gamma_i(t) = 20 \log_{10} \left(\frac{4\pi d_i(t) f_c}{c_0} \right) + \eta_{\text{LoS}}, \quad (6)$$

where f_c is the carrier frequency, c_0 is the speed of light, and η_{LoS} is the additional attenuation factor of the LoS link. The uplink data rate from UAV i to the satellite o at time slot t is

$$r_i(t) = B \log_2 \left(1 + \frac{P_i^{\text{tx}} 10^{-\Gamma_i(t)/10}}{\sigma^2} \right), \quad (7)$$

where B is the uplink channel bandwidth, P_i^{tx} is the transmit power of UAV i , and σ^2 is the power of the additive white Gaussian noise. We denote by r_i the average uplink rate of UAV i over a visibility window.

E. Computation and Communication Cost Model

a) *Computation cost*: Performing K_i queries requires K_i forward evaluations of the denoising loss through the frozen backbone and the adapter. Let C denote the number of CPU cycles required by one forward evaluation, and let f_i be the CPU frequency of UAV i . The computation latency is $T_i^{\text{cmp}} = CK_i/f_i$. The corresponding computation energy is

$$E_i^{\text{cmp}} = \kappa_i CK_i f_i^2, \quad (8)$$

where κ_i is the effective switched capacitance coefficient determined by the chip architecture of UAV i . The computation capability of UAV i is $\theta_i \triangleq 1/(\kappa_i f_i^2)$. Substituting θ_i into (8), the computation energy is rewritten as

$$E_i^{\text{cmp}} = \frac{c_{\text{cmp}}}{\theta_i} K_i, \quad (9)$$

where $c_{\text{cmp}} \triangleq C$ is the per-query computation workload.

b) *Communication cost*: After local computation, UAV i uploads $b_i = K_i/2$ scalars, each of a fixed size s_0 bits, so the uplink payload is $K_i s_0/2$ bits. With the uplink rate r_i , the transmission latency of UAV i is $T_i^{\text{cmm}} = K_i s_0/(2r_i)$. The corresponding transmission energy is

$$E_i^{\text{cmm}} = P_i^{\text{tx}} T_i^{\text{cmm}} = c_{\text{cmm},i} K_i, \quad c_{\text{cmm},i} \triangleq \frac{P_i^{\text{tx}} s_0}{2r_i}, \quad (10)$$

where $c_{\text{cmm},i}$ is the per query communication energy coefficient of UAV i .

F. Satellite Visibility Model

Owing to the orbital motion, the satellite is visible to each UAV only during a fraction of each orbital period. The visibility indicator of UAV i at time slot t is defined as

$$v_i(t) = \begin{cases} 1, & \varphi_i(t) \geq \alpha_e, \\ 0, & \text{otherwise,} \end{cases} \quad (11)$$

where α_e is the minimum elevation threshold. Accordingly, the visibility ratio of UAV i is $\rho_i = \frac{1}{T_{\text{orb}}} \sum_t v_i(t) \Delta t$, where the summation is taken over one orbital period. To guarantee that UAV i can complete and upload its queries within a visibility

window, the total latency must not exceed the per-pass visible duration $\rho_i T_{\text{orb}}$, which can be expressed as

$$T_i^{\text{cmp}} + T_i^{\text{cmm}} \leq \rho_i T_{\text{orb}}. \quad (12)$$

The constraint (12) is equivalent to an upper bound on the number of queries, i.e.,

$$K_i \leq K_i^{\text{max}} \triangleq \left\lfloor \frac{\rho_i T_{\text{orb}}}{C/f_i + s_0/(2r_i)} \right\rfloor, \quad (13)$$

where $\lfloor \cdot \rfloor$ denotes the floor function. Here, K_i^{max} is the maximum number of queries that UAV i can complete within one visibility window, which is determined by its visible duration, channel rate, and known CPU frequency f_i . Therefore, the satellite can obtain K_i^{max} in advance.

III. CONTRACT-BASED INCENTIVE MECHANISM

In the considered system, the satellite aims to obtain more ZO queries with lower payments given to UAVs, whereas each UAV prefers a higher reward with lower energy cost. Due to this objective mismatch and the private computation capability of UAVs, their interaction is formulated as a contract design problem under information asymmetry.

A. UAV's Type and Utility Function

A UAV with capability θ_j is referred to as a type- j UAV, where a larger θ_j indicates a lower per-query computation energy c_{cmp}/θ_j . Since the communication coefficient $c_{\text{cmm},i}$ in (10) and the query cap K_i^{max} in (13) can be obtained by the satellite in advance, they are public information. In contrast, the switched capacitance κ_i depends on the chip architecture and is known only to UAV i , making the computation capability θ_i its private information. Let $\mathcal{J} = \{1, 2, \dots, J\}$ denote the set of computation types, which are ordered as

$$\theta_1 < \theta_2 < \dots < \theta_J. \quad (14)$$

Here, N_j is the number of type- j UAVs, with $\sum_j N_j = I$, and $p_j = N_j/I$ is the fraction of type- j UAVs.

Based on the public communication state, UAVs are divided into groups with the same communication coefficient and query cap, denoted by c_{cmm} and K^{max} , respectively. UAVs with an empty visibility window, i.e., $K_i^{\text{max}} = 0$, are excluded from the current round. Assuming that the computation type is independent of the communication state, the type distribution $p_j * j \in \mathcal{J}$ is preserved within each group. Therefore, we focus on one representative group, for which the satellite designs a type-level contract $(K_j, R_j) * j \in \mathcal{J}$, where K_j and R_j denote the requested number of queries and the reward for a type- j UAV, respectively.

The utility of a type- j UAV that selects the bundle (K_j, R_j) is the reward minus the energy cost of performing and uploading K_j queries, i.e.,

$$U_j = R_j - \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j, \quad (15)$$

where the cost term consists of the computation cost $\frac{c_{\text{cmp}}}{\theta_j} K_j$ in (9) and the communication cost $c_{\text{cmm}} K_j$ in (10).

B. Satellite's Utility Function

The utility of the satellite consists of the query contribution obtained from UAVs and the corresponding rewards paid to them. Specifically, the contribution of type- j UAVs is modeled as a logarithmic function of the number of ZO queries [10]. Using the average per-UAV form, the expected utility is

$$U_o = \sum_{j \in \mathcal{J}} p_j [a \ln(1 + K_j) - R_j], \quad (16)$$

where a is a positive query-value factor and p_j is the fraction of type- j UAVs. The first term $a \ln(1 + K_j)$ represents the value of the contributed ZO queries. As K_j increases, more query evaluations can reduce the variance of the ZO gradient estimate and improve the training quality. However, such improvement gradually saturates, and thus the logarithmic form is adopted to capture the diminishing marginal value of additional ZO queries, which is motivated by the $\mathcal{O}(1/\sqrt{K_j})$ convergence behavior of ZO training [5]. The second term R_j denotes the reward paid to a type- j UAV.

C. Contract Theory and Design Goals

The interactions between the satellite and the UAVs are formulated as a contract theory. The strategy of the satellite is to design the bundles $\{(K_j, R_j)\}_{j \in \mathcal{J}}$ to maximize its utility, while each UAV selects the contract bundle that maximizes its own utility. To induce truthful self-selection, a feasible contract must satisfy the individual rationality (IR) and the incentive compatibility (IC), defined as follows.

Definition 1 (Individual Rationality). *A contract satisfies the IR if each type- j UAV obtains a non-negative utility by selecting the bundle designed for its type, i.e.,*

$$R_j - \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j \geq 0, \quad \forall j \in \mathcal{J}. \quad (17)$$

Definition 2 (Incentive Compatibility). *A contract satisfies the IC if each type- j UAV obtains the highest utility by selecting the bundle (K_j, R_j) designed for its own type rather than any other bundle, i.e.,*

$$R_j - \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j \geq R_{j'} - \left(\frac{c_{\text{cmp}}}{\theta_{j'}} + c_{\text{cmm}} \right) K_{j'}, \quad \forall j, j' \in \mathcal{J}. \quad (18)$$

Under the IC, each UAV is induced to truthfully reveal its private type by self selecting the intended bundle, so that the satellite can screen the heterogeneous UAVs under information asymmetry.

D. Optimization Problem

With the above definitions, the optimal contract problem of the satellite is formulated as

$$\mathbf{P1} : \quad \max_{\{K_j, R_j\}} \quad U_o = \sum_{j \in \mathcal{J}} p_j [a \ln(1 + K_j) - R_j] \quad (19a)$$

$$\text{s.t.} \quad \text{IR constraint in (17),} \quad \forall j \in \mathcal{J}, \quad (19b)$$

$$\text{IC constraint in (18),} \quad \forall j, j' \in \mathcal{J}, \quad (19c)$$

$$K_j \in 2\mathbb{Z}_{\geq 0}, \quad \forall j \in \mathcal{J}, \quad (19d)$$

$$0 \leq K_j \leq K^{\max}, \quad \forall j \in \mathcal{J}. \quad (19e)$$

In **P1**, the decision variables are the query numbers $\{K_j\}$ and the rewards $\{R_j\}$. The constraint (19d) indicates that K_j is an even non-negative integer, since each two-point estimate in (3) consumes two queries, and the constraint (19e) guarantees that every UAV in the group can complete the requested queries within its visibility window.

IV. OPTIMAL CONTRACT ANALYSIS

In this section, we first reduce the IR and IC constraints, derive the optimal reward for any given query sequence, and finally derive the optimal query sequence with an ironing step that ensures the monotonicity.

A. Reduction of Constraints

We first analyze the monotonicity of a feasible contract in the following lemma.

Lemma 1. *For any feasible contract, the requested queries and the rewards are both monotonically non-decreasing in the computation type, i.e., for $j > j'$, if $\theta_j > \theta_{j'}$, then $K_j \geq K_{j'}$ and $R_j \geq R_{j'}$.*

Proof. For any two types $j, j' \in \mathcal{J}$, the IC constraint (18) of type- j against type- j' and that of type- j' against type- j are

$$R_j - \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j \geq R_{j'} - \left(\frac{c_{\text{cmp}}}{\theta_{j'}} + c_{\text{cmm}} \right) K_{j'}, \quad (20)$$

$$R_{j'} - \left(\frac{c_{\text{cmp}}}{\theta_{j'}} + c_{\text{cmm}} \right) K_{j'} \geq R_j - \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j. \quad (21)$$

Adding (20) and (21), the reward terms and the common communication terms $c_{\text{cmm}}(K_j + K_{j'})$ cancel, which yields

$$c_{\text{cmp}} \left(\frac{1}{\theta_{j'}} - \frac{1}{\theta_j} \right) (K_j - K_{j'}) \geq 0. \quad (22)$$

Since $\theta_j > \theta_{j'}$ implies $\frac{1}{\theta_{j'}} - \frac{1}{\theta_j} > 0$ and $c_{\text{cmp}} > 0$, it follows from (22) that $K_j \geq K_{j'}$. Substituting it back into (20) gives $R_j - R_{j'} \geq \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) (K_j - K_{j'}) \geq 0$, i.e., $R_j \geq R_{j'}$. $\square$

Based on Lemma 1, the optimal reward for any given query sequence is derived as follows.

Lemma 2. *Given any monotonic query sequence $\{K_j\}_{j \in \mathcal{J}}$ with $0 \leq K_1 \leq \dots \leq K_J$, the unique optimal reward is*

$$R_j = \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) K_j + \sum_{l=1}^{j-1} c_{\text{cmp}} K_l \left(\frac{1}{\theta_l} - \frac{1}{\theta_{l+1}} \right), \quad \forall j \in \mathcal{J}. \quad (23)$$

Proof. Under the monotonicity in Lemma 1, the global IC reduces to the local downward IC, which binds at the optimum to minimize the reward. Thus, the IR of the lowest type and the local downward IC bind:

$$R_1 = \left(\frac{c_{\text{cmp}}}{\theta_1} + c_{\text{cmm}} \right) K_1, \quad (24)$$

$$R_j = R_{j-1} + \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) (K_j - K_{j-1}), \quad j \geq 2. \quad (25)$$

Unrolling the recursion (25) from (24), the common communication coefficient c_{cmm} cancels in each difference, which

gives (23). Conversely, every term in (23) is non-negative, so the IR holds for all types; together with the monotonicity, the binding local downward IC implies the global IC. $\square$

B. Optimal Query Sequence

By substituting the optimal reward (23) into the satellite's utility (16) and exchanging the order of summation, the expected reward becomes $\sum_j p_j R_j = \sum_j \Lambda_j K_j$, where

$$\Lambda_j = p_j \left(\frac{c_{\text{cmp}}}{\theta_j} + c_{\text{cmm}} \right) + c_{\text{cmp}} \left(\frac{1}{\theta_j} - \frac{1}{\theta_{j+1}} \right) \Omega_j \quad (26)$$

is the virtual marginal cost of type j , with $\Omega_j = \sum_{l=j+1}^J p_l$ and the convention $1/\theta_{J+1} = 0$. In (26), the first term comes from the type- j UAV's cost compensation, and the second term is the information rent surcharge, which aggregates over higher types via Ω_j and arises only from the private computation term c_{cmp} . Substituting the expected reward back, the satellite utility separates across types, and the per type objective is

$$J_j(K_j) = p_j a \ln(1 + K_j) - \Lambda_j K_j. \quad (27)$$

Theorem 1. Suppose that the projected sequence $\{\bar{K}_j\}_{j \in \mathcal{J}}$ in (29) is non-decreasing. Then $\{\bar{K}_j\}_{j \in \mathcal{J}}$ is the optimal query sequence of the continuous relaxation of **P1**, and a feasible near-optimal query sequence of **P1** is obtained by rounding $\{\bar{K}_j\}_{j \in \mathcal{J}}$ monotonically to even integers. Specifically, the unconstrained maximizer is

$$\tilde{K}_j = \frac{p_j a}{\Lambda_j} - 1, \quad (28)$$

its projection onto $[0, K^{\max}]$ is

$$\bar{K}_j = \min \{K^{\max}, \max\{0, \tilde{K}_j\}\}, \quad (29)$$

and the optimal even query number is

$$K_j^* = \arg \max_{K \in \{2\lfloor \bar{K}_j/2 \rfloor, 2\lceil \bar{K}_j/2 \rceil\} \cap [K_{j-1}^*, K^{\max}]} J_j(K), \quad (30)$$

with the convention $K_0^* \triangleq 0$.

Proof. The first-order derivative of J_j in (27) is $\frac{\partial J_j}{\partial K_j} = \frac{p_j a}{1+K_j} - \Lambda_j$, and the second-order derivative is $-\frac{p_j a}{(1+K_j)^2} < 0$, so J_j is concave. Setting the first-order derivative to zero gives the unconstrained maximizer $\tilde{K}_j$ in (28). Since J_j is concave, its maximizer over $[0, K^{\max}]$ is the projection $\bar{K}_j$ in (29), and, by unimodality, its maximizer over the even integers in $[0, K^{\max}]$ is one of the two even integers adjacent to $\bar{K}_j$. Besides, restricting the candidates to $[K_{j-1}^*, K^{\max}]$ in (30) preserves the monotonicity required by Lemma 1, and the candidate set is non-empty since $\{\bar{K}_j\}$ is non-decreasing. $\square$

After obtaining $\{K_j^*\}$, the optimal reward $\{R_j^*\}$ is computed by (23). We note that the projected sequence $\{\bar{K}_j\}$ is not necessarily non-decreasing in practice, because the virtual marginal cost Λ_j in (26) jointly depends on θ_j , θ_{j+1} , the fraction p_j , and Ω_j , so a higher type does not always yield a larger $\bar{K}_j$; meanwhile, since the cap $K^{\max}$ is common to all types in the group, the projection (29) preserves the ordering of $\{\bar{K}_j\}$. To restore the monotonicity required by Lemma 1, a

Algorithm 1 Optimal Contract Design for Satellite-UAV ZO Federated Fine-Tuning

```

1: Input: fractions  $\{p_j\}$ , types  $\{\theta_j\}$ , coefficients  $a$ ,  $c_{\text{cmp}}$ ,  $c_{\text{cmm}}$ , common cap  $K^{\max}$ ;
2: Output: optimal contract  $\{(K_j^*, R_j^*)\}_{j \in \mathcal{J}}$ ;
3: for  $j = J, J-1, \dots, 1$  do
4:   Compute  $\Omega_j$  and  $\Lambda_j$  by (26);
5:   Compute  $\tilde{K}_j$  by (28) and project to  $\bar{K}_j$  by (29);
6: end for
7: while  $\{\bar{K}_j\}_{j \in \mathcal{J}}$  is not non-decreasing do
8:   Find a decreasing sub-sequence and merge it by solving  $\max_{K \in [0, K^{\max}]} \sum_l J_l(K)$ ;
9: end while
10: for  $j = 1, 2, \dots, J$  do
11:   Round  $\bar{K}_j$  to the even integer  $K_j^*$  by (30);
12:   Compute the optimal reward  $R_j^*$  by (23);
13: end for
14: return  $\{(K_j^*, R_j^*)\}_{j \in \mathcal{J}}$ .

```

bunching-and-ironing procedure is adopted, for any decreasing sub-sequence of $\{\bar{K}_j\}_{j \in \mathcal{J}}$, all its elements are merged and assigned a common value by solving the single-variable concave problem $\max_{K \in [0, K^{\max}]} \sum_l J_l(K)$, whose solution admits the closed form $(\sum_l p_l a) / (\sum_l \Lambda_l) - 1$ projected onto $[0, K^{\max}]$, until the whole sequence becomes non-decreasing. The merged values are then rounded by (30). The overall procedure is summarized in Algorithm 1.

V. SIMULATION RESULTS

To evaluate the performance of the proposed scheme, we assume that the system consists of one LEO satellite server and ten UAVs. A conditional U-Net diffusion model with 8.59×10^6 parameters is pretrained on six EuroSAT land-cover classes and frozen [11], while the remaining four classes are learned through a FiLM adapter of 9,472 trainable parameters [4]. The private samples are partitioned by the Dirichlet distribution with parameter 0.5, and the federated ZO fine-tuning lasts 600 rounds with antithetic two point perturbations ($\mu = 5 \times 10^{-3}$) and batch size 64. Each query performs one forward evaluation of 510.3 GFLOP. The onboard accelerators of UAVs have heterogeneous energy efficiencies of $\{1.2, 3, 7.5\}$ TFLOPS/W, corresponding to the $J = 3$ computation types $\theta \in \{0.4, 1, 2.5\}$ with per-query computation energy of 2.1-12.8 mJ. For the uplink, the transmit power is 0.5 W, the bandwidth is 180 kHz, and the median SINR is 0 dB with 2.6 dB log-normal shadowing; each antithetic pair is carried by a 2-kbit packet, yielding per-query communication energy of median 5.6 mJ truncated within $[0.5, 15]$ mJ [12]. Since P1 has no explicit budget term, the budget is met by tuning the query-value factor a : a larger a yields larger K_j^* and hence larger reward, so the expected payment $\sum_j p_j R_j^*$ increases monotonically with a . Unless otherwise specified, all schemes are evaluated under the same per-round budget $\Omega \triangleq \sum_j p_j R_j^* = 20$ (2 J) for fairness. The generation quality is measured by the KID, with




Fig. 2. KID score versus the communication overhead under different schemes.




Fig. 3. Satellite utility versus the satellite incentive budget under different schemes.




Fig. 4. Satellite utility versus the number of UAVs under different schemes.

the Random-Query baseline averaged over three random seeds. We compare the proposed contract with five benchmarks: 1) *Greedy*, allocating queries to the most cost-efficient types; 2) *Bayesian Stackelberg*, optimizing leader rewards under Bayesian type priors; 3) *Equal-Query*, requesting the same queries from all UAVs; 4) *Random-Query*, generating random feasible even allocations; 5) *Fixed-Price*, offering a uniform per-query price.

Fig. 2 compares the uplink communication payload required by different schemes to reach given KID score thresholds. It can be observed that the proposed ZO Contract scheme achieves the same KID targets with lower communication overhead than the baselines. For example, when the target KID score is 75, the ZO-Equal-Query scheme cannot reach this target, while the FO-FedAvg schemes require about 11.4 MB of uplink payload. In contrast, the proposed scheme only requires about 0.2 MB, reducing the communication cost by more than one order of magnitude. These results demonstrate that the proposed scheme can effectively improve communication efficiency in resource constrained satellite UAV scenarios.

Fig. 3 shows the satellite utility versus the budget Ω . It can be observed that the utility of each scheme grows with Ω at a diminishing rate, owing to the logarithmic satisfaction. The proposed contract always achieves the highest utility, since it induces each UAV to truthfully reveal its private computation type and assigns more queries to the cost-efficient UAVs, thereby extracting more queries per unit payment. Fig. 4 shows the satellite utility versus the number of UAVs. The utility of all schemes increases with more participants, and the performance gain of the proposed contract gradually enlarges, which demonstrates the good scalability of the proposed scheme.

VI. CONCLUSION

In this paper, we have proposed a game-theoretic ZO federated fine-tuning framework for GenAI in satellite-UAV networks. Specifically, we have developed a satellite-UAV federated diffusion learning model, where the LEO satellite coordinates the UAVs to fine-tune the diffusion model through a lightweight adapter and ZO optimization with only scalar uplinks. Then, we have formulated the interactions between the satellite and the UAVs as a query-reward contract design problem under information asymmetry, and derived the optimal contract in closed form with the monotone rounding and

bunching-and-ironing steps. Simulation results have demonstrated that the proposed scheme achieves a higher satellite utility while maintaining model quality with substantially lower uplink cost.

REFERENCES

- [1] Y. Tan, T. Koketsu Rodrigues, N. Kato, M. Ariyoshi, and Y. Hasegawa, "Pivotal ai paradigms for satellite communication security: From machine learning to large language models," *IEEE Communications Surveys Tutorials*, vol. 28, pp. 4654–4689, 2026.
- [2] X. Tang, Q. Chen, W. Weng, C. Jin, Z. Liu, J. Wang, G. Sun, X. Li, and D. Niyato, "Task assignment and exploration optimization for low altitude uav rescue via generative ai enhanced multi-agent reinforcement learning," *IEEE Transactions on Mobile Computing*, vol. 25, no. 1, pp. 627–643, 2026.
- [3] P. Peng, X. Chen, T. Xu, C. Wu, Y. Zou, Q. Ni, and E. Soljanin, "Uncertain location transmitter and uav-aided warden-based leo satellite covert communication systems," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 11 099–11 112, 2026.
- [4] A. Shysheya, J. Bronskill, M. Patacchiola, S. Nowozin, and R. E. Turner, "FiT: Parameter efficient few-shot transfer learning for personalized and federated image classification," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023.
- [5] M. Assaad, Z. Nehme, and M. Debbah, "Communication-efficient zero-order and first-order federated learning methods over wireless networks," *IEEE Transactions on Signal Processing*, pp. 1–16, 2026.
- [6] L. Xie, Z. Su, Y. Wang, and Z. Li, "A practical federated learning framework with truthful incentive in uav-assisted crowdsensing," *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 129–144, 2025.
- [7] J. Ru, R. Wang, and M.-S. Alouini, "Coverage and rate analysis of follower-based leo satellite networks: A stochastic geometry approach," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 12 662–12 675, 2026.
- [8] Y. Guo, Y. Chen, H. Li, Y. Wu, and J. Huang, "Aoi-aware incentive mechanism for uav-assisted mobile crowdsensing: A contract-theoretic approach," *IEEE Transactions on Mobile Computing*, vol. 25, no. 2, pp. 1660–1677, 2026.
- [9] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [10] J. Huang, B. Ma, Y. Wu, Y. Chen, and X. Shen, "A hierarchical incentive mechanism for federated learning," *IEEE Transactions on Mobile Computing*, vol. 23, no. 12, pp. 12 731–12 747, 2024.
- [11] P. Helber, B. Bischke, A. Dengel, and D. Borth, "EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
- [12] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, "Energy efficient federated learning over wireless communication networks," *IEEE Transactions on Wireless Communications*, vol. 20, no. 3, pp. 1935–1949, 2021.