

Special Topic: Integration of Large AI Model and 6G

Satellite edge artificial intelligence with large models: architectures and technologies[†]

Yuanming SHI¹, Jingyang ZHU^{1*}, Chunxiao JIANG²,
Linling KUANG² & Khaled Ben LETAIEF³¹*School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China*²*Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 100084, China*³*Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong 999077, China*

Received 31 October 2024/Revised 21 February 2025/Accepted 30 April 2025/Published online 20 June 2025

Abstract Driven by the growing demand for intelligent remote sensing (RS) applications, large artificial intelligence (AI) models (LAMs) pretrained on large-scale unlabeled datasets and fine-tuned for downstream tasks have considerably improved learning performance for various downstream tasks because of their generalization capabilities. However, many specific downstream tasks, such as extreme weather nowcasting (e.g., downbursts and tornadoes), disaster monitoring, and battlefield surveillance, require real-time data processing. Traditional methods that transfer raw data to ground stations for processing often cause substantial issues in terms of latency and trustworthiness. To address these challenges, satellite edge AI provides a paradigm shift from ground-based to on-board data processing by leveraging the integrated communication and computation capabilities in space computing power networks, thereby enhancing the timeliness, effectiveness, and trustworthiness for RS downstream tasks. Moreover, satellite edge LAMs involve the training (i.e., fine-tuning) and inference phases, where a key challenge lies in developing computation task decomposition principles to support scalable LAM deployment in resource-constrained space networks with time-varying topologies. In this article, we first propose a satellite federated fine-tuning architecture to split and deploy the modules of LAM over space and ground networks for efficient LAM fine-tuning. Then, we introduce a microservice-empowered satellite edge LAM inference architecture that virtualizes LAM components into lightweight microservices tailored for multitask multimodal inference. Finally, we discuss the future directions for enhancing the efficiency and scalability of satellite edge LAM, including task-oriented communication, brain-inspired computing, and satellite edge AI network optimization.

Keywords large AI models (LAMs), satellite edge LAM, space-computing power networks (Space-CPNs), satellite federated fine-tuning, microservice-empowered satellite edge LAM inference

Citation Shi Y M, Zhu J Y, Jiang C X, et al. Satellite edge artificial intelligence with large models: architectures and technologies. *Sci China Inf Sci*, 2025, 68(7): 170302, <https://doi.org/10.1007/s11432-024-4425-y>

1 Introduction

Remote sensing (RS) satellite Earth observation systems have global coverage, all-weather, and full-spectrum detection capabilities and play a crucial role in global economic development, resource management, and environmental protection. Multi-resolution, multi-sensor, and multi-temporal unlabeled RS data can provide different temporal-spatial-spectral descriptive information of objects, empowering various downstream tasks, such as environmental monitoring, extreme weather forecasting, disaster monitoring, and military surveillance [1]. Artificial intelligence (AI) methods, represented by deep learning, have considerably improved the efficiency of RS data interpretation. Currently, an important development direction in the field of intelligent RS data interpretation is to leverage large-scale multimodal data, enabling AI models to have multitask generalization capabilities while achieving efficient and high-precision inference. The approach of pretraining and fine-tuning large AI models (LAMs) for downstream tasks is a proven effective solution. RS large models, which are pretrained on large-scale unlabeled datasets (e.g., fMoW [2] and BigEarthNet [3]) using self-supervised learning techniques, have achieved substantial improvement in performance. Existing LAMs have adopted various neural network architectures,

* Corresponding author (email: zhuji2@shanghaitech.edu.cn)

[†] Special Topic: The 27th Annual Meeting of the China Association for Science and Technology

including diffusion models (e.g., DiffusionSat [4]) and Transformer-based models (e.g., SatMAE [5] and SpectralGPT [6]).

Many specific RS downstream tasks have high requirements for data timeliness. For instance, in the field of extreme weather forecasting, severe convective weather events, such as downbursts, tornadoes, and hailstorms, are characterized by sudden onset, strong locality, short lifespan, and remarkable impact, making them a persistent challenge in weather forecasting (i.e., on the scale of hours) and nowcasting (i.e., on the scale of minutes) tasks [7]. Directly downloading such massive RS data for ground processing would inevitably lead to severe latency issues and pose risks of privacy breaches [8]. For instance, in 2022 alone, Landsat-8 from the Landsat program of the National Aeronautics and Space Administration generated 73.8 TB of publicly available RS imagery over regions in China. Meanwhile, the transmission rates of common feeder links typically range from several hundred Mbps to a few Gbps [9]. Therefore, the development of more efficient architectures for scalable and low-latency intelligent data processing becomes critical. To address this substantial demand for computing power, space computing power networks (Space-CPNs) have emerged as a promising solution to provide global connectivity and support ubiquitous intelligent services, encompassing aircraft and various satellites in space. The on-board processing capabilities of satellites have considerably improved, especially with respect to graphics processing units (GPUs) designed for neural network training. The GPU for the Space HPC program, supported by the European Space Agency and led by the Barcelona Supercomputing Center, successfully validated the feasibility, low power consumption, and high performance of integrating GPUs on satellites in 2019 [10]. Moreover, the Φ -Sat-1 mission of the European Space Agency marked the first successful use of deep neural networks to assist Earth observation tasks [11]. In this context, complex RS downstream services at the satellite edge can be enhanced by leveraging LAMs, which benefit from the continuously advancing computational and communication capabilities of nodes within Space-CPNs. The overall architecture of satellite edge AI is illustrated in Figure 1. Satellite edge AI with LAMs primarily involves two major types of computation tasks: training and inference.

Training tasks in satellite edge LAM refer to the fine-tuning of the LAMs using the satellite sensing data to enhance the timeliness and trustworthiness for the downstream tasks. This process is constrained by limited computing power at the satellite edge, making it difficult to meet the high computation demands of large model fine-tuning tasks, even with the collaboration among multiple satellites. Therefore, the exploration of the computation task decomposition approaches of satellite fine-tuning tasks for LAMs and the development of joint communication and computation orchestration methods are key to addressing the issue of insufficient single-point computing power and supporting on-board fine-tuning [12, 13]. In this article, we propose a satellite federated fine-tuning architecture, establishing an intelligent coordination mechanism between Space-CPN elements and ground facilities to enhance the timeliness of satellite fine-tuning tasks and improve the efficiency of network resource utilization. In particular, the LAM modules are partitioned, with the parameter-heavy components deployed on ground cloud servers, while the satellites utilize their computing payloads to fine-tune the classification layers of LAM. This requires designing new inter-satellite scheduling schemes and satellite-ground coordinated transmission technologies to support the satellite federated fine-tuning.

Satellite edge LAM inference refers to the provision of on-board interpretation empowered by LAMs for various RS downstream applications requested from ground users, thereby improving the overall inference timeliness and accuracy [14]. These multimodal inference tasks involve the sequential processing of requested data through the main modules of LAMs, including modality encoders, input projectors, a backbone calculator, an output projector, and a modality decoder. Many downstream inference tasks require the use of the same multiple modality encoders and input projectors, resulting in substantial computation demands and increased latency, which leads to remarkable computation redundancy, where identical modality encoders and mappers are repeatedly scheduled across different inference tasks. Furthermore, in traditional multimodal multitask inference, where different modules are consolidated into a single multimodal network for a specific inference task, common encoding and projection modules are often redundantly deployed for different tasks [15]. To address these issues, we propose a microservice-empowered satellite edge LAM inference architecture for multitask multimodal inference. In this architecture, the functional modules (e.g., modality encoders and input projectors) of LAM are virtualized into a series of independent microservices, presenting the advantages of low coupling and high cohesion. However, because of the time-varying property of satellite network topology and the heterogeneity of computation resources, the efficient deployment of microservice modules and their sequential orchestration at the satellite edge remain critical challenges. This article discusses intuitive methods for microservice

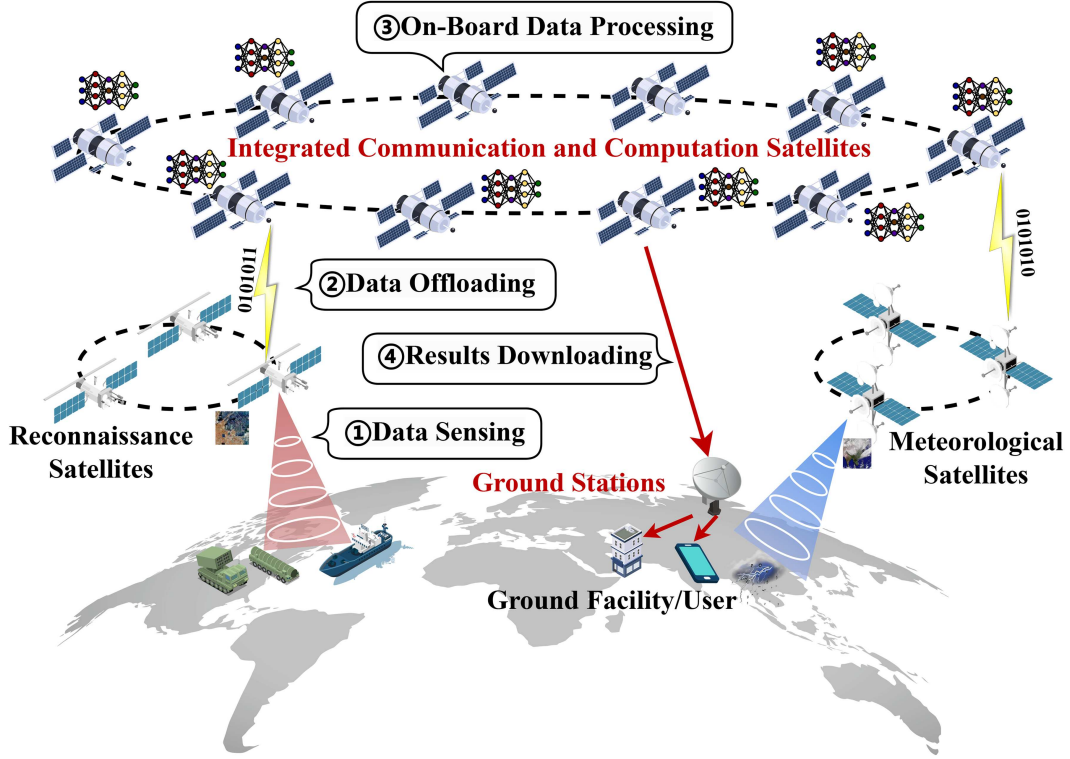


Figure 1 (Color online) Procedures of satellite edge artificial intelligence, including data sensing, data offloading, on-board data processing, and results downloading.

deployment and orchestration at the satellite edge to address these challenges.

The main contributions of this article are three-fold. First, we propose a satellite federated fine-tuning architecture tailored for satellite edge training by partitioning and distributing the modules of LAMs and developing corresponding inter-satellite link (ISL) scheduling and satellite-ground coordinated transmission schemes to support this architecture in Section 2. To support inference service for various multimodal downstream tasks, we introduce a microservice-empowered satellite edge LAM inference architecture to address the problem of computation redundancy originating from multitask multimodal LAM inference, followed by providing solutions for the microservice deployment and orchestration problems over Space-CPNs in Section 3. Finally, we outline our future work from three perspectives: task-oriented communications, neuromorphic computing, and generative AI for optimizing satellite edge AI networks, in Section 4.

2 Satellite federated fine-tuning for LAM training

In this section, we present the satellite federated fine-tuning architecture, the ISL scheduling schemes, and the satellite-ground coordinated transmission techniques.

2.1 Satellite federated fine-tuning architecture

With the deployment of numerous RS satellite projects, such as the Landsat series, Gaofen series, and Sentinel-2, satellites equipped with various types of sensors can capture massive multimodal data, including multispectral, hyperspectral, infrared, and synthetic aperture radar images. These RS data can support a wide range of Earth observation applications, including land use classification, weather forecasting, and disaster monitoring. The success of RS large models is demonstrated by their significant enhancement of the feature representation capabilities of traditional data-driven models and their improved generalization performance on downstream tasks. In this article, we utilize a state-of-the-art RS large model, i.e., SpectralGPT [6]. SpectralGPT adopts a vision Transformer (ViT) [16] as its network backbone and is pretrained on several large-scale multispectral datasets comprising millions of images with varying sizes, time series information, and Earth regions. The ViT consists of an embedding layer

for image patching and position embedding, several transformer encoders, and a multi-layer perceptron (MLP) head for classification. Typically, ViT has three versions with different parameter sizes, i.e., ViT-base with 86 million parameters, ViT-large with 307 million parameters, and ViT-huge with 632 million parameters. Specifically, a masked autoencoder [17], an emerging baseline method, is used in the pretraining phase of SpectralGPT, where unmasked patches or pixels are employed to reconstruct the masked patches or pixels in a self-supervised manner. As a result, in the pretraining phase, an encoder for learning visual representations and a decoder for image reconstruction can be obtained, where the decoder replaces the MLP head during pretraining and the MLP head is used for fine-tuning downstream tasks. Then, the RS large models can be applied to other downstream tasks by fine-tuning them on specific datasets, which aims at enhancing the learning performance by updating the parameters of the LAM, instead of training from the beginning [18].

A traditional method of pretraining and fine-tuning LAMs using massive space data is downloading them to the ground facilities for centralized training. However, downloading raw satellite data violates some data privacy protocols and suffers from large communication overhead. For instance, Sentinel-2 of the European Space Agency, as part of its Earth observation satellite constellation, is expected to generate 1.6 TB of high-resolution optical imagery data within each observation cycle across its various orbits. To this end, we propose a satellite federated fine-tuning scheme to update local models on-board without transmitting raw data to ground stations (GSs). However, fine-tuning the full parameter LAM (e.g., SpectralGPT with ViT-huge backbone that has 632 million parameters) on-board results in high computation and communication complexity and storage pressure for resource-constrained satellites. Parameter-efficient fine-tuning methods, such as prompt tuning, prefix tuning, adapter tuning, and low-rank adapters, have been proposed to reduce the number of fine-tuning parameters for pretrained LAM while achieving comparable learning performance with fine-tuning the full parameter LAM [19]. Instead, we propose to adopt parameter-efficient head tuning, where the parameters of the encoder in the foundation model are kept frozen and only the head layer is fine-tuned for different downstream tasks. For example, in single-label and multi-label classification tasks, the original MLP can be fine-tuned; meanwhile, for semantic segmentation or change detection tasks, UPerNet layers can be added after the MLP layer for head tuning [6]. Given the limited processing capabilities of satellites and abundant cloud resources on the ground, we propose a unique model deployment scheme tailored for satellite federated head tuning, where the embedding and head layers of the foundation model SpectralGPT [6] are deployed on low-Earth-orbit (LEO) satellites, while the encoder containing most of the model parameters is deployed on ground cloud servers. To elaborate, the embedding and head layers of SpectralGPT have approximately 50000 and 62000 parameters, respectively, accounting for only 0.13% of the total parameters in ViT-base. In this setup, the satellite federated fine-tuning task aims at training the local head layer on each LEO satellite using their local data and obtaining a global head after model aggregation.

The detailed communication and computation process of the proposed satellite federated fine-tuning architecture is illustrated in Figure 2, which consists of two phases, i.e., feature extraction in Figure 2(b) and model update in Figure 2(c). In the feature extraction phase, each LEO satellite first computes embedding vectors by transforming local data into patches and then inputs them into the embedding layer. Subsequently, one LEO satellite within each orbit gathers the embedding vectors from all satellites via intra-orbit ISLs and sends the concatenated intra-orbit embedding vectors to the GSs via satellite-ground links (SGLs). GSs relay the concatenated embedding vectors to the ground cloud server via ground-dedicated links. Then, the ground cloud server extracts feature vectors from the embedding vectors via the deployed encoder and sends them back. Once the LEO satellites receive the feature vectors, the model update phase begins. The LEO satellites first train their local heads and then perform an intra-orbit model aggregation via ISLs to accelerate the model fine-tuning. After several intra-orbit communication rounds, each satellite in one orbit obtains the same intra-orbit head and delivers the updated head to the GSs for global aggregation via SGL. Finally, a global head is broadcast to all satellites for training in the subsequent communication round. This process repeats until the fine-tuning procedure converges.

2.2 ISL scheduling and satellite-ground coordinated transmission

Specifically, in the satellite federated fine-tuning architecture, although it is unnecessary to transmit terabytes of data to the ground cloud server, several challenges remain to be addressed. First, although the embedding layer placed on the LEO satellite has a relatively small number of parameters, the output

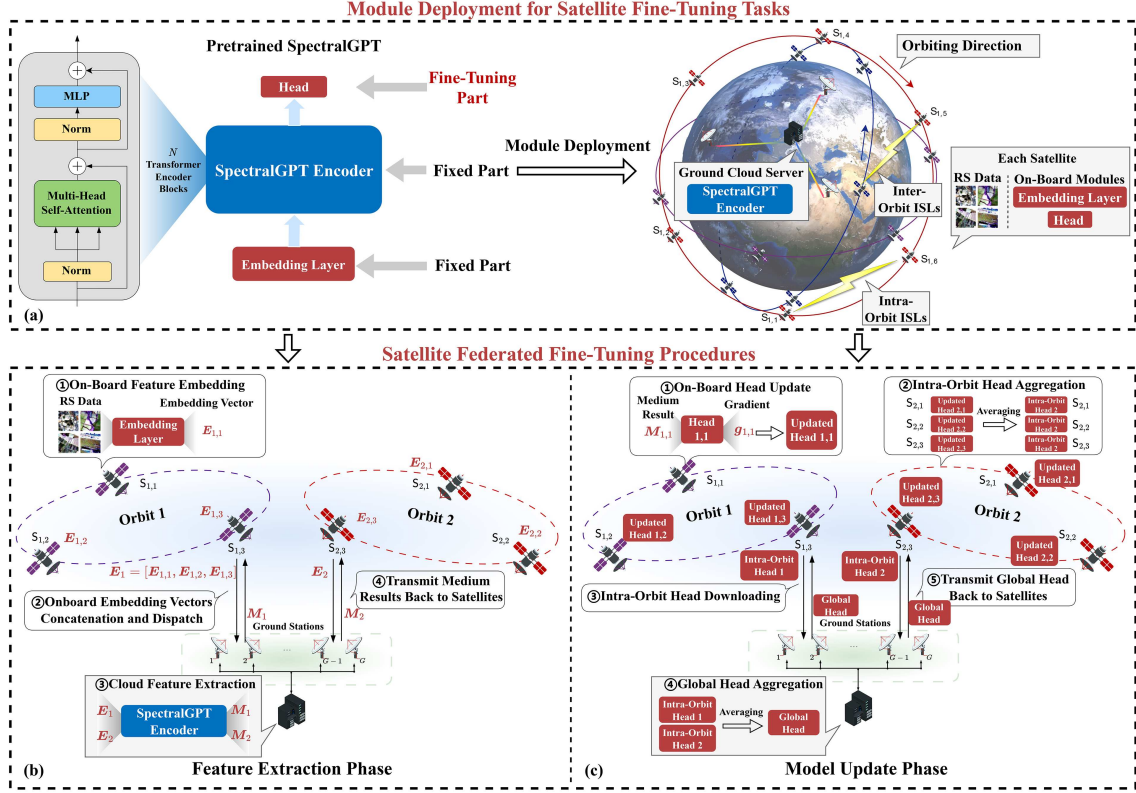


Figure 2 (Color online) Satellite federated fine-tuning architecture, including LAM module deployment and operating procedures. (a) Module deployment for satellite federated fine-tuning task; (b) feature extraction phase; (c) model update phase.

embedding vectors can be substantial because of the large number of high-resolution images resized and input into the embedding layer of SpectralGPT. For example, using the EuroSAT dataset with a batch size of 512, the output embedding vector dimension is 786432, and when stored with a 32-bit float precision, each batch consumes approximately 1.5 GB of memory. Given that the Ka/Ku band feeder links between LEO satellites and GSs generally have capacities ranging from several hundred megabits per second to a few gigabits per second and the SGL durations typically last only a few minutes, downloading all embedding vectors from satellites in an orbit to the ground still presents significant challenges [9]. Second, with the development of laser ISL technology [20], LEO satellites can transfer data between one another at speeds of several to tens of gigabits per second via laser ISLs. Effectively leveraging inter-satellite communication during the feature extraction phase for transmitting embedding vectors and utilizing intra-orbit aggregation during the model update phase to accelerate the convergence of fine-tuning tasks are crucial. In summary, new task-oriented inter-satellite and satellite-ground transmission approaches need to be proposed to support satellite federated fine-tuning tasks.

2.2.1 ISL scheduling schemes

During the satellite federated fine-tuning process, massive data are exchanged through intra-orbit ISLs, including embedding vectors and local heads. To improve the transmission efficiency and accelerate model fine-tuning, we propose a ring all-reduce-based algorithm to support both intra-orbit embedding gathering and model aggregation, which leverages the stable ring topology of LEO satellites in one orbit. First, the on-board data are divided into several blocks. Then, each satellite transmits a certain block to its neighbor in parallel. At the end of transmission, each satellite has the complete data of two certain blocks. To proceed, each satellite transmits the local complete data block to its neighborhood until each satellite possesses information from all satellites in orbit. With this parallel transmission scheme, the transmission time is no longer related to the number of LEO satellites in one orbit but only to the size of the data. Notably, the rate of ISLs may vary between different satellites or change over time in real systems, and the overall performance is constrained by the slowest link. In addition, ISLs between adjacent satellites within the same orbit may be affected by alignment losses, leading to potential transmission outages [21].

The global model aggregation with the ground cloud server as the central coordinator is limited by the low data rates of SGLs and single-point failure issues, meaning that if the ground cloud server fails, then the training process cannot continue. To solve this problem, the decentralized head tuning can also be considered with complete on-board model aggregation via intra-orbit and inter-orbit ISLs [20]. Existing decentralized designs on the ground cannot be directly utilized in LEO satellite networks because of the special topology of satellite networks. Specifically, high-speed and stable intra-orbit ISLs with the ring topology guarantee fast intra-orbit transmission [22]. Furthermore, inter-orbit ISLs are time-varying and unreliable. For instance, in several constellations, such as Kepler, OneWeb, and Starlink, the inter-orbit ISL rates primarily range from 1 to 4 Gbps [23], which is significantly lower than that of intra-orbit ISLs. ISLs established between satellites orbiting in opposite directions are referred to as cross-seam inter-orbit ISLs [24], which suffer from Doppler shift (reaching up to 2.4 GHz) and short duration because of the rapid movement of the satellites, resulting in an increased bit error rate during inter-orbit data transmission. Therefore, the use of inter-orbit links needs to be strategically scheduled to maintain the overall efficiency and reliability of the system. To accelerate the inter-orbit transmission process, a set of paths is selected to transmit the model in parallel by setting the weight of each ISL as the reciprocal of the link capacity. Subsequently, we propose to utilize the Floyd-Warshall algorithm to find the shortest path between each pair of satellites [25]. The shortest path between the first orbit and the last orbit is selected, and any path that overlaps the shortest path is removed to avoid congestion. This process is repeated, and the selected paths are used to transmit data in parallel. The tailored inter-orbit communication scheme solves the single-point failure problem and fully utilizes satellite transmission resources to transmit models in parallel to minimize latency.

2.2.2 *Satellite-ground coordinated transmission schemes*

Through multi-beam and multiple access technologies, a single LEO satellite can simultaneously connect to multiple GSs, and a GS can establish SGLs with multiple satellites within its line of sight. Given that all satellites within the same orbit have an identical intra-orbit model upon completing intra-orbit aggregation, the model from each orbit can be transmitted collaboratively by several LEO satellites to multiple GSs, significantly reducing the data load required to be transmitted over each SGL. Specifically, the satellite-ground communication topology in each communication window can be modeled as a static flow network, which contains edges, satellite vertices, GS vertices, a sink vertex, and a virtual source vertex. The weight of the edge between the source vertex and a satellite vertex is determined by the fraction of untransmitted data in the corresponding orbit during the current communication window. The weight of the edge between a satellite vertex and a GS vertex is determined by the maximum proportion of data that can be transferred via the corresponding link during the current communication window. The capacity of the edge between a GS vertex and the sink vertex is determined by the maximum proportion of data that can be transferred via the dedicated link during the current communication window. This problem of maximizing the overall data rate during satellite-ground data transmission can be formulated as a classical network flow problem and solved by adopting the well-known Ford-Fulkerson algorithm [26]. The network flow model simulates the flow of water through pipes to identify the maximum data transmission path in a capacity-constrained network, ensuring that data are transmitted from satellites to GSs in the most efficient manner. The network flow problem is illustrated in Figure 3.

Because of the rapid development and deployment of LEO satellite array antennas, SGL is increasingly adopting multiple-input multiple-output (MIMO) communication technology [27]. In addition, over-the-air computation (AirComp) emerges as a promising approach to achieve a more efficient aggregation process. AirComp is designed to compute a class of nomographic functions over wireless channels by leveraging the waveform superposition property of multiple access channels [28]. With this technique, signals from each device are transmitted simultaneously, enabling the summation operation to be performed directly over the air. A potential application of AirComp could involve coordinated downlink transmission between LEO satellite swarms and GSs, which ensures notable spectral efficiency improvements because of both the increased spatial separation of antennas and the inherent benefits of AirComp. However, the high mobility of LEO satellites presents new challenges for AirComp-based MIMO SGL systems, particularly in acquiring accurate channel state information (CSI) within the limited coherence time. To mitigate these challenges, the predictable movement patterns of LEO satellites can be used to approximate CSI by considering the positional relationship between transmit and receive antennas, along with the estimated atmospheric influences. This approach enables the integration of communica-

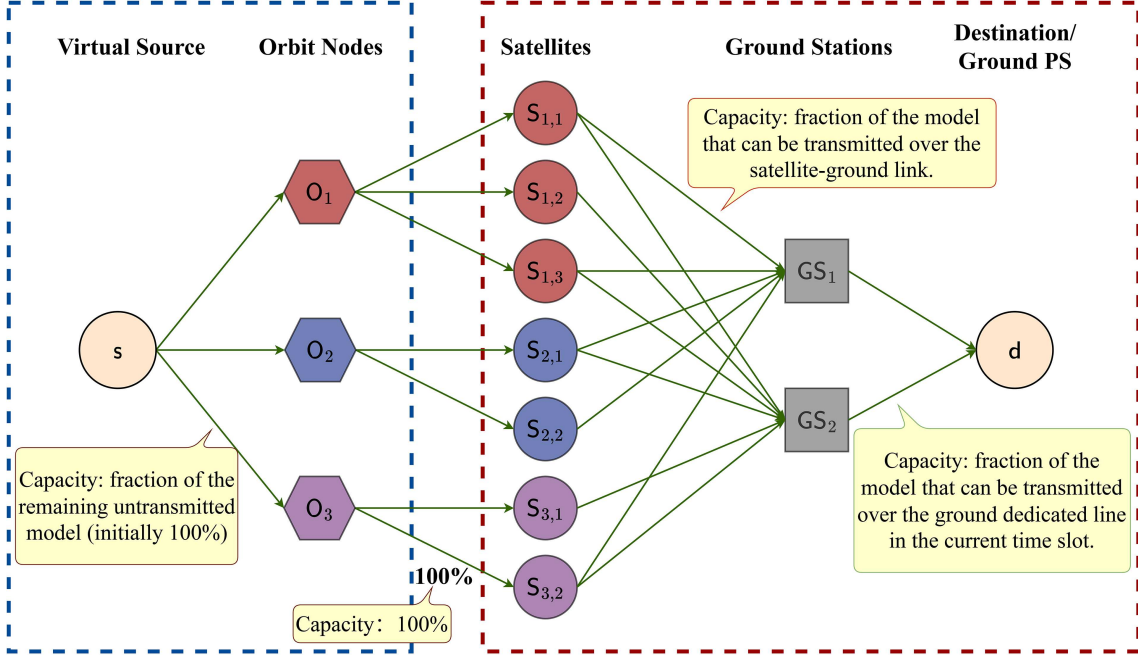


Figure 3 (Color online) Formulated network flow problem for satellite-ground coordinated transmission.

tion and computation, ultimately enhancing the spectral efficiency and reducing the aggregation latency of SGL systems. Nevertheless, affected by various factors, such as strong interference, unstable weather conditions, and significant atmospheric attenuation, the communication link between GSs and LEO satellites typically experiences a low signal-to-noise ratio, which inevitably leads to severe distortion if analog-modulation-based AirComp is used to aggregate data in satellites. To address this issue, lattice quantization provides an effective solution for enabling coded AirComp, ensuring strong noise resistance while maintaining compatibility with existing digital systems [29]. Moreover, the lattice coding approach for AirComp eliminates the need for CSI at the transmitter, which makes it particularly suitable for SGL systems. In conclusion, the lattice coding approach for AirComp integrates the benefits of both digital and analog transmission, as discussed in [30]. This method enhances bandwidth efficiency, improves system robustness against channel conditions, and reduces latency.

3 Microservice-empowered satellite edge inference for LAM

In this section, we present the microservice-empowered satellite edge inference architecture for multitask multimodal LAM inference and the developed microservice deployment and orchestration schemes over Space-CPN.

3.1 Microservice-empowered satellite edge LAM inference architecture

Mainstream multimodal LAMs comprise modality encoders, input projectors, a backbone calculator, an output projector, and a modality decoder, facilitating various inference tasks through the sequential processing of requested data (e.g., text and image modal data) via these modules. For example, applications, such as generating videos from text and images, removing clutter from images based on textual prompts, or intelligent interpretation of RS data, all fall under the category of multimodal LAM inference tasks. Multitask inference in multimodal LAMs necessitates the activation of multiple modality encoders and input projectors, resulting in substantial computation demands and increased latency. Existing approaches mitigate computation latency by minimizing the calculation overhead of the overall compressed model, which ignores the heightened response latency caused by the repeated scheduling of identical modality encoders and mappers across different inference tasks [31]. Furthermore, in traditional multimodal multitask inference, which consolidates different modules into a single multimodal network for a specific inference task, common encoding and projection modules are often redundantly deployed across various networks.

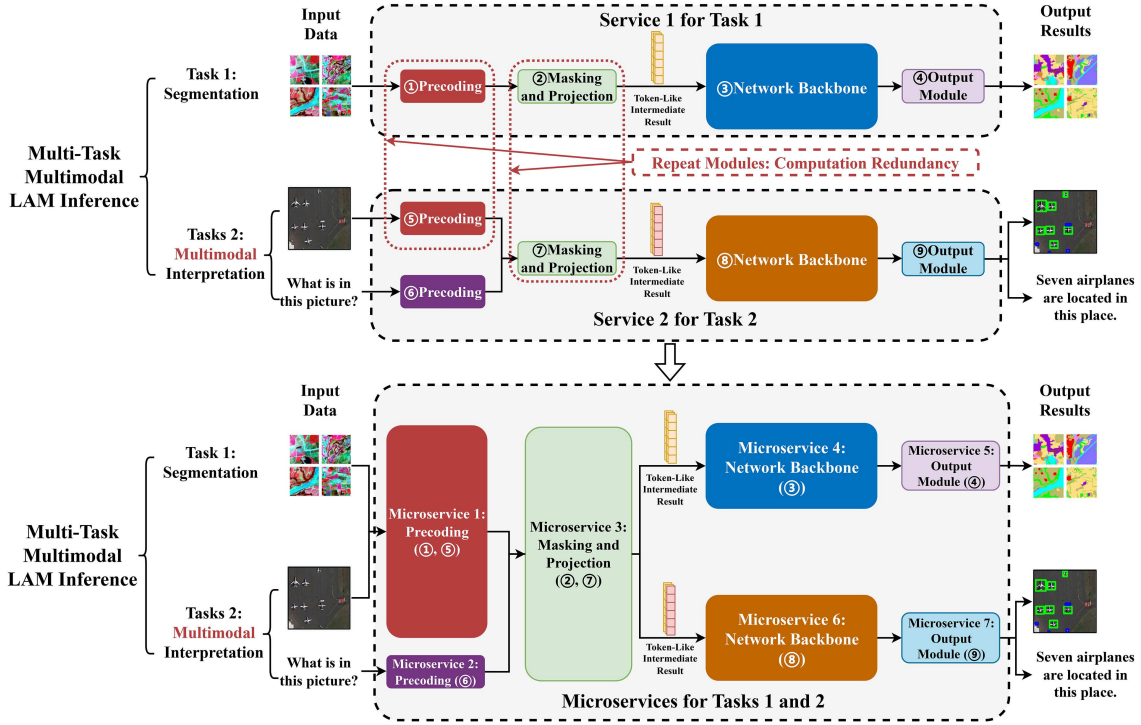


Figure 4 (Color online) Microservice-empowered satellite edge LAM inference architecture. In the presented scenario, precoding, masking, and projection modules for Tasks 1 and 2 are redundant. As a result, Services 1 and 2 can be decomposed into several sequential microservices, with Microservices 1 and 3 serving as shared modules for Tasks 1 and 2, thereby addressing computation redundancy.

To address the computation redundancy problem and further reduce inference latency, we propose a microservice-empowered satellite edge LAM inference architecture for multiple multimodal downstream tasks, which virtualizes the functional modules of the multimodal LAM into a series of independent microservices [32]. As illustrated in Figure 4, traditional service-based multitask multimodal LAM inference packages the entire computation process into a single service, requiring a complete invocation for each inference task, which overlooks the shared modules between multiple tasks. Specifically, the microservice-empowered satellite edge LAM inference architecture transforms the classical monolithic inference service, typically deployed in data centers, into a unidirectional acyclic graph of distinct microservices, which can be distributed across networks. As shown in Figure 4, for Tasks 1 and 2, Microservices 1 and 3 can be simultaneously invoked for the shared modules, and the remaining microservice modules in the original computation sequence can be utilized to complete the multitask multimodal satellite edge LAM inference. This architecture has the advantages of agility, portability, scalability, and resilience in software development and maintenance. Given that frequent packing and unpacking operations over the intermediate activations are performed for exchanges among microservices, the extra communication overhead is optimized by a latency-aware granularity of the microservice scheduling approach while improving resource utilization.

As illustrated in Figure 5, various microservice modules need to be first pre-deployed on each satellite node. For each inference request from the ground user involving RS data sensing, intelligent inference, and result transmission, the complex computation task is decomposed into multiple submodules or microservice modules, which are assigned to the corresponding satellite nodes containing the necessary modules for sequential execution. Then, the final results are transmitted back to the ground user. This satellite edge inference process has two main challenges. The first challenge is the “where to compute” issue concerning microservice deployment [14]. Given a known satellite routing table, the goal is to deploy the appropriate microservice modules on each satellite based on the statistical requirements of the inference tasks. The second challenge is the “how to compute” issue related to microservice orchestration [33]. With the deployed microservice modules, this issue involves scheduling the computing resources of satellite nodes according to the execution sequence of microservice modules required by real-time inference tasks to complete the computation process.

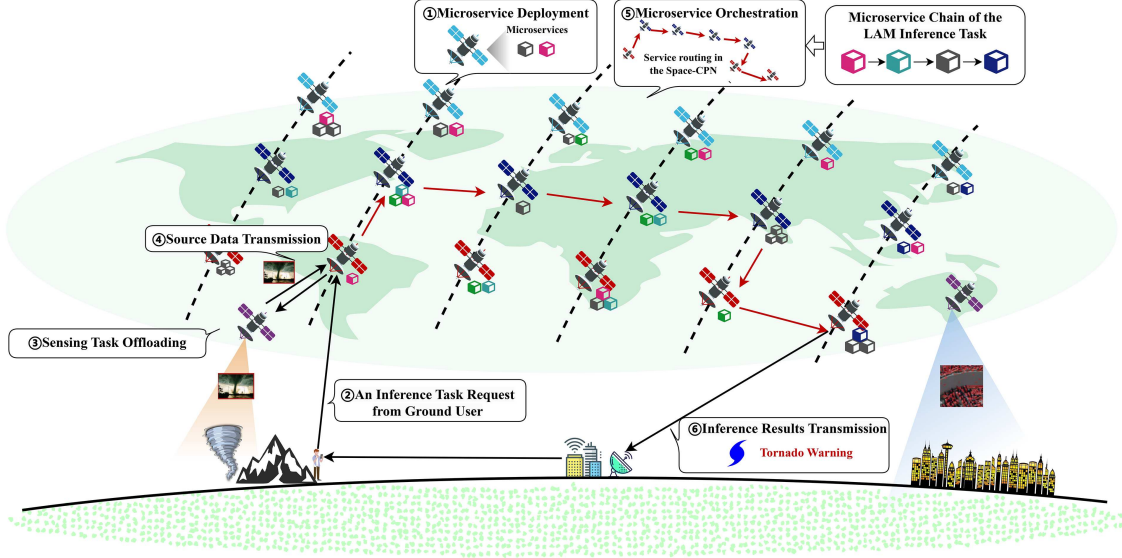


Figure 5 (Color online) Procedures of microservice deployment and orchestration for microservice-empowered satellite edge inference with LAMs, including microservice deployment, inference request, sensing task offloading, source data transmission, microservice orchestration, and inference results transmission.

3.2 Microservice deployment and orchestration

Space-CPN provides global coverage, and because of the significant differences in terrain and population density across various regions, the demands for RS and computation services vary considerably. For example, satellites over vast ocean areas may have abundant idle computing resources, while satellites over densely populated or high-priority monitoring areas face relatively constrained computing resources. Spatially, the task load on satellites varies significantly across different regions. Temporally, the geographical position of each satellite changes rapidly, and as it passes over regions with different user densities and RS needs, its task load also fluctuates quickly. In contrast to terrestrial networks, idle computing resources in satellite networks are distributed unevenly across both spatial and temporal dimensions. Furthermore, the heterogeneity of satellite computing and communication capabilities (e.g., CPUs, memory, and energy) in Space-CPN, the uncertainty of the dynamic topology of the satellite network, and the interdependencies between different microservice modules make it crucial to efficiently plan and orchestrate computing resources within the network to effectively complete inference tasks.

3.2.1 Microservice deployment

In this article, microservice deployment in LEO satellite networks refers to the strategic allocation of computation and communication tasks involved in executing multitask multimodal LAM inference at the satellite edge [14]. Although the performance of on-board multitask multimodal LAM inference has benefits, such as reduced data transmission to Earth and faster response times, the intricate microservice deployment strategies and dynamic satellite topologies have significant challenges. The limited computing power, energy constraints, and unpredictable communication links of LEO satellites further complicate the management of microservices, particularly when handling diverse and heterogeneous inference tasks. These inference tasks can be modeled as directed acyclic graphs (DAGs), where nodes represent different microservices and edges represent the data dependencies between them. Specifically, the latency in a microservice-empowered inference architecture consists of processing, propagation, and transmission latencies, whereas the total execution latency for a given task is determined by the slowest path among all microservices involved. We suppose that multiple inference tasks need to be executed, each of which consists of multiple microservices operating sequentially. The goal of microservice deployment is to optimize the deployment of these microservices within the satellite networks to minimize the total execution latency of all inference tasks while efficiently utilizing the limited resources of heterogeneous satellites. The goal of microservice deployment is to minimize the total latency while efficiently utilizing the limited resources of heterogeneous satellites. The problem can be modeled as a matching between the DAG of the inference tasks and the time-varying satellite networks for latency minimization. The

resulting optimization problem is a mixed-integer linear programming (MILP) problem because of the integer constraints of the deployment indicator, which is an NP-hard problem.

To solve the MILP problem considering the time-varying nature of the transmission and propagation latencies of ISLs, traditional solvers cannot be directly applied because some objective parameters, such as transmission latency, are functions of optimization variables that are unknown during the modeling process. The problem is further reformulated as a Markov decision process (MDP), where the states of the satellite networks, including available resources and communication conditions, are considered. Moreover, the actions represent deployment choices, and the rewards are based on the latency and resource efficiency of each decision. We further propose a reinforcement learning (RL)-based solution using proximal policy optimization (PPO), which enables the system to learn how to deploy microservices in a manner that reduces latency and avoids resource bottlenecks. The use of PPO ensures an adaptable and scalable solution that effectively handles the complex dependencies of microservices across a constellation of LEO satellites. The resulting deployment strategies are designed to minimize latency and resource consumption, ensuring efficient execution of multitask multimodal LAM inference in a resource-constrained and highly dynamic environment. Notably, the microservice deployment problem in this article can be extended to a multi-agent RL approach [34]. Specifically, we can partition the satellite networks into several subnetworks, where each subnetwork consists of multiple satellites and is managed by a single agent. The agents communicate with each other to exchange local subnetwork information and adjust their strategies accordingly while sharing the same reward function to ensure cooperative decision-making across the satellite networks.

3.2.2 *Microservice orchestration*

In this article, we define microservice orchestration in the satellite edge inference architecture as the coordination of computation and communication tasks incurred by designing a service routing scheme for terrestrial requests. Microservice orchestration aims at finding an energy-efficient microservice routing scheme for these requests over large-scale mega-constellations with hundreds to thousands of satellites because of the bottleneck on satellites. Specifically, the microservices for satellite edge inference have been deployed on the satellites in advance. We suppose that at least two or more microservices collaborate to respond to the request and model the requested service as a DAG to represent the execution processes of microservices. Then, we construct an augmented directed graph, where the weight of each edge represents the energy consumption of both computation and communication when orchestrating microservices over the satellite networks. If each satellite in the network is equipped with all possible microservice modules, then the microservice orchestration problem will be reduced to a minimum spanning tree problem. However, because it is impractical for each satellite to host all microservice modules, relay nodes must be introduced. As a result, we can reformulate the problem of energy-efficient microservice routing as a minimum directed Steiner tree (DST) problem to minimize the total energy consumption of satellite edge inference [35]. We further propose a heuristic algorithm to approximate the optimal solution by first using the Dijkstra algorithm to find the minimum distance paths and then merging the paths. By solving the DST problem, the resulting microservice orchestration policies can help minimize the total energy consumption of the microservice processing and data transmission procedures for microservice-empowered satellite edge inference.

Although heuristic and approximating algorithms for solving the DST problem are suitable for static topologies, they may perform poorly when applied to highly dynamic networks with time-varying topologies, such as satellite networks. Thus, solutions for the DST problem under time-varying topologies need to be explored. Traditional neural networks struggle to capture the structural information of nodes and edges in a graph, whereas graph neural networks (GNNs), a type of neural network specifically designed to process graph data, can effectively extract useful structural features through the message passing mechanism between neighboring nodes. In addition, RL can incrementally build partial solutions by learning a policy to make optimized choices during the decision-making process without requiring pre-labeled data. Consequently, combining RL with GNNs becomes an ideal choice for solving the DST problem in the presence of time-varying topologies [36]. In the DST problem, RL is modeled as an MDP, with the following elements: state space is the state of the current partial solution on the graph, including selected nodes and edges, graph information, and problem information. Action space is the selection of a new node to add to the current solution in the partial solution. The reward function is the reward at each step equals the negative weight of the selected edge. Policy selects actions based on the estimated Q value

function, where the Q value represents the expected future cumulative reward after taking a particular action. For the choice of deep Q -network, the graph attention network (GAT) can be employed because it can capture the relationships between nodes in a graph and update the embedding of each node through message passing mechanisms. In the solving process, we input the graph-structured data, i.e., the state, into the GAT, and the output is the action for network decision-making.

4 Future directions

In this section, we discuss future directions for satellite edge LAM, encompassing three aspects, i.e., task-oriented communications for downstream tasks, more energy-efficient computing architectures, and the application of generative AI in optimizing satellite edge AI networks.

4.1 Task-oriented communications for satellite edge LAM

Task-oriented communication provides a solution for satellite edge AI inference by extracting key features and eliminating redundant information, thereby improving the efficiency of information transmission between satellites and GSs [37]. In terms of satellite edge AI, satellites are equipped with various sensors that capture multimodal data, such as optical imagery, infrared, and radar signals. In contrast to traditional multimodal data, RS multimodal data exhibit distinct characteristics, often being multispectral and capturing observations from various angles. The resolution of data collected by different sensors can vary significantly. Satellites can also observe the same location over extended periods, providing a valuable temporal dimension. These multimodal RS data are capable of providing high-quality features for downstream tasks, such as change detection and cloud removal. However, existing data interpretation methods suffer from several limitations. First, traditional methods require downloading raw data to GSs for post-processing, which is at odds with the limited spectral resources of SGLs. Second, traditional single-sensor, data-driven, unimodal deep learning methods struggle to achieve optimal performance when applied to different modalities. Furthermore, traditional multimodal approaches often fail to effectively align data based on the characteristics of each modality, making it difficult to fuse data from different imaging mechanisms and extract meaningful insights. Finally, meteorological effects, such as atmospheric turbulence, pose significant challenges for the data sensing, processing, and transmission procedures in SGLs [38]. To address these challenges, we propose a solution based on the multimodal information bottleneck (MIB) for RS downstream tasks that aims at reducing the volume of data transmitted through SGLs while maximizing the extraction of features from multimodal RS data. The information bottleneck [39] is originally designed as a task-oriented communication method that effectively characterizes the trade-off between informativeness and compression. This approach aims at extracting features that are highly informative for downstream tasks while minimizing task-irrelevant information, thereby reducing transmission overhead.

The MIB approach provides a robust framework for this purpose, enabling dimensionality reduction within each feature space to minimize information loss and alignment of heterogeneous output features to ensure more effective fusion [40]. First, the MIB approach identifies the compressed representation for each input modality that maximizes the retention of information relevant to the target task while minimizing redundancy. Then, these compressed representations are fused into an efficient and robust encoded feature vector. The GSs receive this noise-corrupted feature vector to derive the inference result for downstream tasks. The MIB approach focuses on extracting the most relevant features from each modality that are crucial for a specific task. By emphasizing task-relevant features, the model effectively reduces the dimensionality of feature vectors, which is particularly advantageous in bandwidth-constrained satellite communication environments. In SGLs, noise can significantly distort the transmitted signal. We introduce the variational information bottleneck approach for maintaining high-quality inference capabilities, even when the transmitted feature vectors are significantly corrupted by noisy environments. This approach improves the reliability of data transmission and stabilizes communication, even in the face of environmental disturbances [41].

4.2 Neuromorphic computing for energy-efficient satellite edge LAM

Energy consumption is becoming the main bottleneck for LAM-empowered computation tasks, and the situation is even worse for satellites, which have limited energy-harvesting capabilities and constrained

battery lifespan [42]. As a result, a more energy-efficient neural network model for LAMs is essential. To this end, a spiking neural network (SNN), a third-generation neural network, has exhibited great potential for LAMs with its unique spike-based event-driven mechanism [43]. In contrast to traditional neural networks, SNN generally uses a leaky integrate-and-fire model to characterize neurons, where a gradually decreased membrane potential is adopted to record the temporal dynamics. The weighted sum of the output of the previous layer contributes to this membrane potential during forward propagation, and each neuron emits a spike as output if its membrane potential reaches a threshold. Based on this special neuron structure, the computation flow can be described as the transition of sparse, discrete, and single-bit spikes. By mimicking the behavior of human neurons with extremely sparse computation, the architecture of LAMs can be redesigned to a more energy-efficient version. In particular, in [44], a spike-driven transformer architecture is proposed for the basic block of the large language models (LLMs). As for vision generation tasks where diffusion models are widely employed, SNN is also attractive because of its biological plausibility compared with traditional artificial neural networks [45]. Overall, these existing architectures confirm that SNN is a promising solution for energy-efficient LAMs.

Another way to address the energy consumption issue is to develop innovative hardware computing architectures. Traditionally, AI models are mainly deployed and executed by advanced modern microprocessors, such as CPUs and GPUs, which have a strong computation capability for complex AI tasks. To fully unleash the potential of these aforementioned SNN-based architectures in energy efficiency, neuromorphic processors are a necessary part that acts as hardware for model deployment [43]. Different from the classical von Neumann architecture, where processing and storage capabilities are separately enabled by CPUs and memory units, the computation task is completed via membrane potential dynamics on neurons and asynchronous binary spikes on synapses, which effectively avoids the throughput limitation caused by the von Neumann bottleneck. In summary, brain-inspired SNNs substantially reduce computational complexity and enhance energy efficiency because of their binarized input and output. In addition, their event-driven nature enables the explicit encoding of input and output into time-dimensioned spike sequences for processing, effectively extracting temporal features, which makes neuromorphic computing well-suited for satellite edge AI for on-board environmental monitoring tasks [46]. The Neuro SatCom project of the European Space Agency has already evaluated the employment of neuromorphic computing to on-board applications. In particular, for an interference detection task, AI models on the non-neuromorphic platform Xilinx Versal consume 136.9 J per problem, while AI models on the neuromorphic platform Spinnaker consume only 6.3 J per problem. Therefore, neuromorphic computing is particularly suitable for resource-limited satellites, which can be further designed to support LAMs [47].

4.3 Generative AI for optimizing satellite edge AI networks

The allocation of communication, computation, storage, and sensing resources plays a critical role in the efficiency of satellite edge AI networks [48]. To enhance the efficiency of satellite edge AI systems, finding effective solutions for large-scale resource optimization problems is of paramount importance. The development of generative AI provides new approaches to solving optimization problems in satellite edge AI systems. For instance, the microservice deployment problem discussed in Subsection 3.2.1, which can be formulated as a MILP problem, can be solved using a diffusion model [49]. First, the process formulates the original problem as an energy minimization task. During the forward diffusion phase, noise is added to an initial near-valid solution, transforming it into a more randomized form, which enables the exploration of the solution space. Then, the reverse diffusion phase removes noise in small steps to reconstruct a solution close to the optimum, driven by minimizing the reverse Kullback-Leibler divergence. GNNs are also integrated to capture complex relationships in graph-structured data, improving solution quality. Annealing techniques are used to gradually reduce temperature, reducing randomness and focusing on low-energy regions for convergence toward a global optimum. As another example, solving the DST problem under time-varying topologies, as discussed in Subsection 3.2.2, can also be approached using generative AI methods. Specifically, by leveraging the capability of the transformer to capture the relevance of historical information, the model can better utilize the correlation between satellite topologies at adjacent time steps, selecting actions to optimize the reward in the RL-based solution to the DST problem [50]. The core challenge of this approach is how to encode graph-structured data into semantically meaningful vector representations for inputting into the transformer. To address this issue, the encoding principles of a graph transformer can be used to encode graph-structured data [51], thereby replacing the GAT-based deep Q-network with a transformer architecture.

Although these methods have achieved certain success in solving their respective optimization problems in satellite edge AI systems, they are limited to addressing the current problem, i.e., the specific data and environment distribution, and cannot be generalized to solve other problems. Therefore, we attempt to explore a “one model for all problems” learning model, utilizing LAMs to provide a unified solution framework for large-scale satellite AI system optimization problems. The authors in [52] proposed NetLLM, an LLM adaptation framework for networking, which is a specific category of LAMs. NetLLM integrates a multimodal encoder for networking data and introduces networking-specific heads for task-specific outputs. The data-driven low-rank networking adaptation reduces fine-tuning costs by adding lightweight, low-rank matrices without altering the pretrained LLM. NetLLM effectively handles different tasks, such as viewport prediction, adaptive bitrate streaming, and cluster job scheduling. To this end, LAMs specifically designed for satellite AI system optimization are highly promising for microservice deployment, orchestration, and migration. Moreover, LLMs can be considered solvers for large-scale optimization problems. LLMs present a promising avenue for democratizing optimization by enabling individuals to describe optimization tasks in natural language, reducing the need for specialized expertise. However, LLMs face challenges in mathematical optimization, including limited context windows that restrict the detailed representation of large-scale problems with multiple constraints, and difficulty in navigating non-convex or highly irregular search spaces [53]. To address these limitations, recent studies in the academic community have begun exploring various strategies. In [54], the authors proposed an LLM-based agent framework called OptiMUS to simplify the use of domain-specific tools, such as Gurobi, through text prompts. Similarly, the authors in [55] proposed an OptLLM framework that employs multi-round dialogues to iteratively refine the modeling and solving of optimization problems. In contrast to the aforementioned works that trained an LLM for one class or a few classes of MILP problems, the authors in [56] proposed to train a single foundation model on diverse MILP classes and instances, thereby improving the generalization capability to unseen classes. In summary, fully exploring and harnessing the potential of LLMs as solvers for satellite AI system optimization problems is another significant direction for future research.

5 Conclusion

Satellite edge AI represents a new paradigm for space data processing, shifting the focus from ground-based to on-board processing techniques. This paradigm involves disruptive technologies for the fine-tuning and inference of LAMs at the satellite edge. To enhance the timeliness, effectiveness, and trustworthiness of RS downstream tasks for LAM-empowered satellite edge AI, the development of computation task decomposition principles is the key to supporting scalable LAM deployment in resource-constrained Space-CPNs with dynamic topology. This article proposes a satellite federated fine-tuning architecture for fine-tuning LAMs by partitioning and distributing LAM modules and developing corresponding ISL scheduling and satellite-ground coordinated transmission schemes. In addition, we introduced a microservice-empowered satellite edge LAM inference architecture to address the computational redundancy problem in multitask multimodal LAM inference and provided insightful solutions for microservice deployment and orchestration problems at the satellite edge. We outlined our future work from three perspectives: task-oriented communications, neuromorphic computing, and generative AI for optimizing satellite edge AI networks. We hope this article will inspire more innovative architectures, technologies, and applications of LAMs for satellite edge AI.

Acknowledgements This work was supported in part by National Natural Science Foundation of China (Grant Nos. 62271318, 62325108, 62341131) and Hong Kong Research Grants Council under the Areas of Excellence Scheme (Grant No. AoE/E-601/22-R).

References

- 1 Sun X, Tian Y, Lu W X, et al. From single- to multi-modal remote sensing imagery interpretation: a survey and taxonomy. *Sci China Inf Sci*, 2023, 66: 140301
- 2 Christie G, Fendley N, Wilson J, et al. Functional map of the world. In: *Proceedings of IEEE Conference Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, 2018. 6172–6180
- 3 Sumbul G, Charfuelan M, Demir B, et al. BigEarthNet: a large-scale benchmark archive for remote sensing image understanding. In: *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Yokohama, 2019. 5901–5904
- 4 Khanna S, Liu P, Zhou L, et al. DiffusionSat: a generative foundation model for satellite imagery. In: *Proceedings of the International Conference on Learning Representations (ICLR)*, Vienna, 2024
- 5 Cong Y, Khanna S, Meng C, et al. SatMAE: pre-training transformers for temporal and multi-spectral satellite imagery. In: *Proceedings of Neural Information Processing Systems (NeurIPS)*, New Orleans, 2022. 197–211
- 6 Hong D, Zhang B, Li X, et al. SpectralGPT: spectral remote sensing foundation model. *IEEE Trans Pattern Anal Mach Intell*, 2024, 46: 5227–5244

- 7 Lebedev V, Ivashkin V, Rudenko I, et al. Precipitation nowcasting with satellite imagery. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, 2019. 2680–2688
- 8 Coffey M M. Balancing privacy rights and the production of high-quality satellite imagery. *Environ Sci Technol*, 2020, 54: 6453–6455
- 9 Horst Y, Bitachon B I, Kulmer L, et al. Tbit/s line-rate satellite feeder links enabled by coherent modulation and full-adaptive optics. *Light Sci Appl*, 2023, 12: 153
- 10 Kosmidis L, Rodriguez I, Jover Á, et al. GPU4S: embedded GPUs in space—latest project updates. *Microprocessors MicroSyst*, 2020, 77: 103143
- 11 Giuffrida G, Fanucci L, Meoni G, et al. The Φ -Sat-1 mission: the first on-board deep neural network demonstrator for satellite earth observation. *IEEE Trans Geosci Remote Sens*, 2022, 60: 1–14
- 12 Lin Z, Qu G, Chen X, et al. Split learning in 6G edge networks. *IEEE Wireless Commun*, 2024, 31: 170–176
- 13 You X H, Huang Y M, Zhang C, et al. When AI meets sustainable 6G. *Sci China Inf Sci*, 2025, 68: 110301
- 14 Yu Z, Cheng H V, Yang Z, et al. Microservice deployment for satellite edge AI inference via deep reinforcement learning. In: Proceedings of the 35th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Valencia, 2024. 1–6
- 15 Al-Doghman F, Moustafa N, Khalil I, et al. AI-enabled secure microservices in edge computing: opportunities and challenges. *IEEE Trans Serv Comput*, 2023, 16: 1485–1504
- 16 Dosovitskiy A, Beyer L, Kolesnikov A, et al. An image is worth 16×16 words: Transformers for image recognition at scale. In: Proceedings of the International Conference on Learning Representations (ICLR), 2021
- 17 He K, Chen X, Xie S, et al. Masked autoencoders are scalable vision learners. In: Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, 2022. 16000–16009
- 18 Dong Z, Gu Y, Liu T. UPetu: a unified parameter-efficient fine-tuning framework for remote sensing foundation model. *IEEE Trans Geosci Remote Sens*, 2024, 62: 1–13
- 19 Jiao L, Huang Z, Lu X, et al. Brain-inspired remote sensing foundation models and open problems: a comprehensive survey. *IEEE J Sel Top Appl Earth Obs Remote Sens*, 2023, 16: 10084–10120
- 20 Wang G, Yang F, Song J, et al. Free space optical communication for inter-satellite link: architecture, potentials and trends. *IEEE Commun Mag*, 2024, 62: 110–116
- 21 Nie S, Akyildiz I F. Channel modeling and analysis of inter-small-satellite links in terahertz band space networks. *IEEE Trans Commun*, 2021, 69: 8585–8599
- 22 Wu S P, Liao W. Scalable contact graph modeling for low earth orbit satellite mega-constellations. *IEEE Netw*, 2025, 39: 257–262
- 23 Lagunas E, Chatzinotas S, An K, et al. *Non-geostationary Satellite Communications Systems*. London: IET, 2022
- 24 Leyva-Mayorga I, Soret B, Popovski P. Inter-plane inter-satellite connectivity in dense LEO constellations. *IEEE Trans Wireless Commun*, 2021, 20: 3430–3443
- 25 Gallo G, Pallottino S. Shortest path algorithms. *Ann Oper Res*, 1988, 13: 1–79
- 26 Shi Y, Zeng L, Zhu J, et al. Satellite federated edge learning: architecture design and convergence analysis. *IEEE Trans Wireless Commun*, 2024, 23: 15212–15229
- 27 Heo J, Sung S, Lee H, et al. MIMO satellite communication systems: a survey from the PHY layer perspective. *IEEE Commun Surv Tut*, 2023, 25: 1543–1570
- 28 Wang Z, Zhao Y, Zhou Y, et al. Over-the-air computation for 6G: foundations, technologies, and applications. *IEEE Int Things J*, 2024, 11: 24634–24658
- 29 Azimi-Abarghouei S M, Varshney L R. Federated learning via lattice joint source-channel coding. In: Proceedings of IEEE International Symposium on Information Theory (ISIT), Athens, 2024. 1623–1628
- 30 Yao J, Xu W, Yang Z, et al. Wireless federated learning over resource-constrained networks: digital versus analog transmissions. *IEEE Trans Wireless Commun*, 2024, 23: 14020–14036
- 31 Xu M, Niyato D, Zhang H, et al. Joint foundation model caching and inference of generative ai services for edge intelligence. In: Proceedings of IEEE Conference on Global Communications (GLOBECOM), Kuala Lumpur, 2023. 3548–3553
- 32 Yang Z, Yu Z, Liu X, et al. Latency-aware microservice deployment for edge AI enabled video analytics. In: Proceedings of IEEE Wireless Communications and Networking Conference (WCNC), Dubai, 2024. 1–6
- 33 Guo H, Cao H, He J, et al. POBO: safe and optimal resource management for cloud microservices. *SIGMETRICS Perform Eval Rev*, 2024, 51: 20–21
- 34 Xu W, Yang Z, Ng D W K, et al. Edge learning for B5G networks with distributed signal processing: semantic communication, edge computing, and wireless sensing. *IEEE J Sel Top Signal Process*, 2023, 17: 9–39
- 35 Charikar M, Chekuri C, Cheung T, et al. Approximation algorithms for directed steiner problems. *J Algorithm*, 1999, 33: 73–91
- 36 Yen B P C, Luo Y. Deep reinforcement learning for solving directed steiner tree problems. In: Proceedings of IEEE Region 10 Conference (TENCON), Hong Kong, 2022. 1–5
- 37 Shi Y, Zhou Y, Wen D, et al. Task-oriented communications for 6G: vision, principles, and technologies. *IEEE Wireless Commun*, 2023, 30: 78–85
- 38 Bui V P, Dinh T Q, Leyva-Mayorga I, et al. Semantic image encoding and communication for earth observation with LEO satellites. *IEEE Trans Cogn Commun Netw*, 2025, 11: 1210–1224
- 39 Tishby N, Pereira F C, Bialek W. The information bottleneck method. In: Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing, 1999. 368–377
- 40 Mai S, Zeng Y, Hu H. Multimodal information bottleneck: learning minimal sufficient unimodal and multimodal representations. *IEEE Trans Multimedia*, 2022, 25: 4121–4134
- 41 Xie S, Ma S, Ding M, et al. Robust information bottleneck for task-oriented communication with digital modulation. *IEEE J Sel Areas Commun*, 2023, 41: 2577–2591
- 42 Yang Y, Xu M, Wang D, et al. Towards energy-efficient routing in satellite networks. *IEEE J Sel Areas Commun*, 2016, 34: 3869–3886
- 43 Schuman C D, Kulkarni S R, Parsa M, et al. Opportunities for neuromorphic computing algorithms and applications. *Nat Comput Sci*, 2022, 2: 10–19
- 44 Yao M, Hu J, Zhou Z, et al. Spike-driven Transformer. In: Proceedings of Neural Information Processing Systems (NeurIPS), New Orleans, 2023. 64043–64058
- 45 Cao J, Wang Z, Guo H, et al. Spiking denoising diffusion probabilistic models. In: Proceedings of IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Waikoloa, 2024. 4912–4921
- 46 Bose P, Kasabov N K, Bruzzone L, et al. Spiking neural networks for crop yield estimation based on spatiotemporal analysis of image time series. *IEEE Trans Geosci Remote Sens*, 2016, 54: 6563–6573
- 47 Lagunas E, Ortiz F G, Eappen G, et al. Performance evaluation of neuromorphic hardware for onboard satellite communication applications. *IEEE Wireless Commun*, 2024, 31: 78–84
- 48 Sheng M, Zhou D, Bai W G, et al. Coverage enhancement for 6G satellite-terrestrial integrated networks: performance metrics, constellation configuration and resource allocation. *Sci China Inf Sci*, 2023, 66: 130303

- 49 Sanokowski S, Hochreiter S, Lehner S. A diffusion model framework for unsupervised neural combinatorial optimization. In: Proceedings of the 41st International Conference on Machine Learning, Vienna, 2024
- 50 Hu S, Shen L, Zhang Y, et al. On transforming reinforcement learning with transformers: the development trajectory. *IEEE Trans Pattern Anal Mach Intell*, 2024, 46: 8580–8599
- 51 Ying C, Cai T, Luo S, et al. Do Transformers really perform badly for graph representation? In: Proceedings of Neural Information Processing Systems (NeurIPS), 2021. 28877–28888
- 52 Wu D, Wang X, Qiao Y, et al. NetLLM: adapting large language models for networking. In: Proceedings of the ACM SIGCOMM 2024 Conference, Sydney, 2024. 661–678
- 53 Yang C, Wang X, Lu Y, et al. Large language models as optimizers. In: Proceedings of the International Conference on Learning Representations (ICLR), Vienna, 2024
- 54 AhmadiTeshnizi A, Gao W, Udell M. OptiMUS: scalable optimization modeling with (MI)LP solvers and large language models. In: Proceedings of the 41st International Conference on Machine Learning, Vienna, 2024
- 55 Zhang J, Wang W, Guo S, et al. Solving general natural-language-description optimization problems with large language models. In: Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), Mexico City, 2024. 483–490
- 56 Li S, Kulkarni J, Menache I, et al. Towards foundation models for mixed integer linear programming. 2024. ArXiv:2410.08288