The 22nd International Conference on Pervasive Computing and Communications (PerCom 2024)

"Stitching Satellites to the Edge: Pervasive and Efficient Federated LEO Satellite Learning"

March 13, 2024, Biarritz, France

Mohamed Elmahallawy and Tie Luo

Computer Science Department, Missouri University of Science and Technology, USA

MISSOURI S&T | 150

Low Earth Orbit (LEO) Satellite Networks

The **recent surge** of interest and investment in large-scale LEO satellites

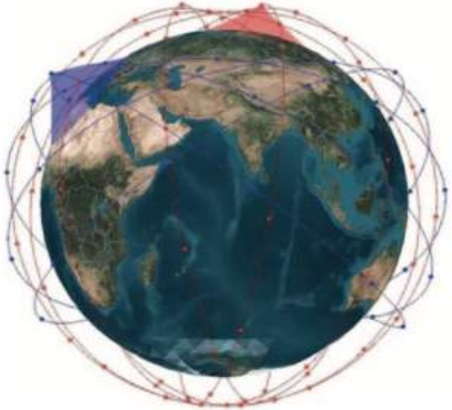

Industry

SPACEX
STARLINK

amazon
Project Kuiper

eutelsat
ONEWEB

Government

NASA

esa
European Space Agency

Low Earth Orbit (LEO) Satellite Networks

Enable novel applications empowered by machine learning, such as:

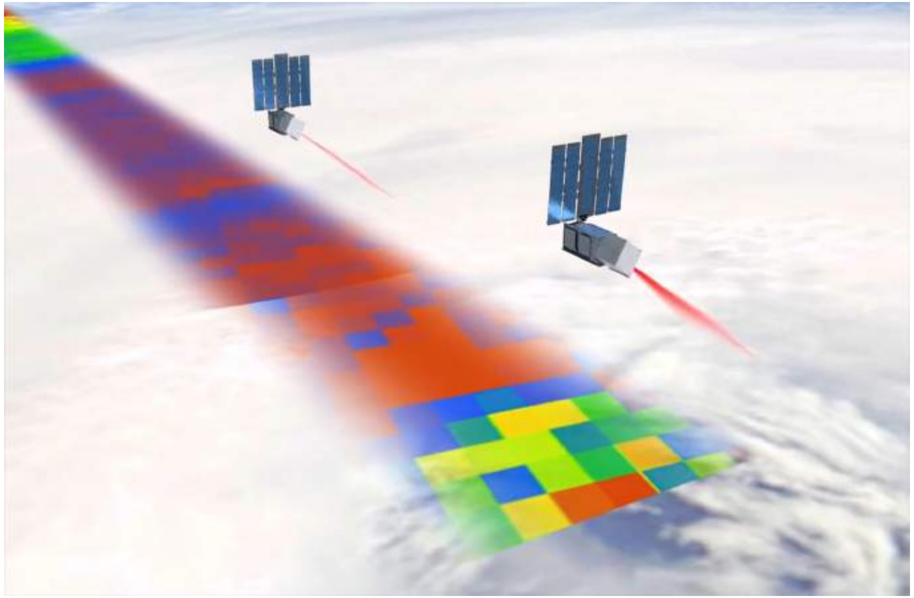

5 petabytes of image data per day (2019)!

Disaster Detection

Border Monitoring

[1] P. Wang, H. Li, B. Chen, and S. Zhang, "Enhancing earth observation throughput using inter-satellite communication," IEEE Transactions on Wireless Communications, vol. 21, no. 10, pp. 7990–8006, 2022

Low Earth Orbit Satellite Networks

Conventional (i.e., centralized) ML:

- Download high-resolution satellite images to a ground station (GS)
- This is impractical because:
  - **Bandwidth**
    - 50~500MB (vs. 5 petabytes satellite data!)
  - **Privacy**
    - Raw data transmission

Consequence of delay

Four–hour delay in downloading satellite images resulted in the **burning** of **135,000 acres** and the loss of **50 lives**[1].

[1] Campfire in Berkeley, California, in November 2018

Introducing Federated Learning (FL) into LEO Networks

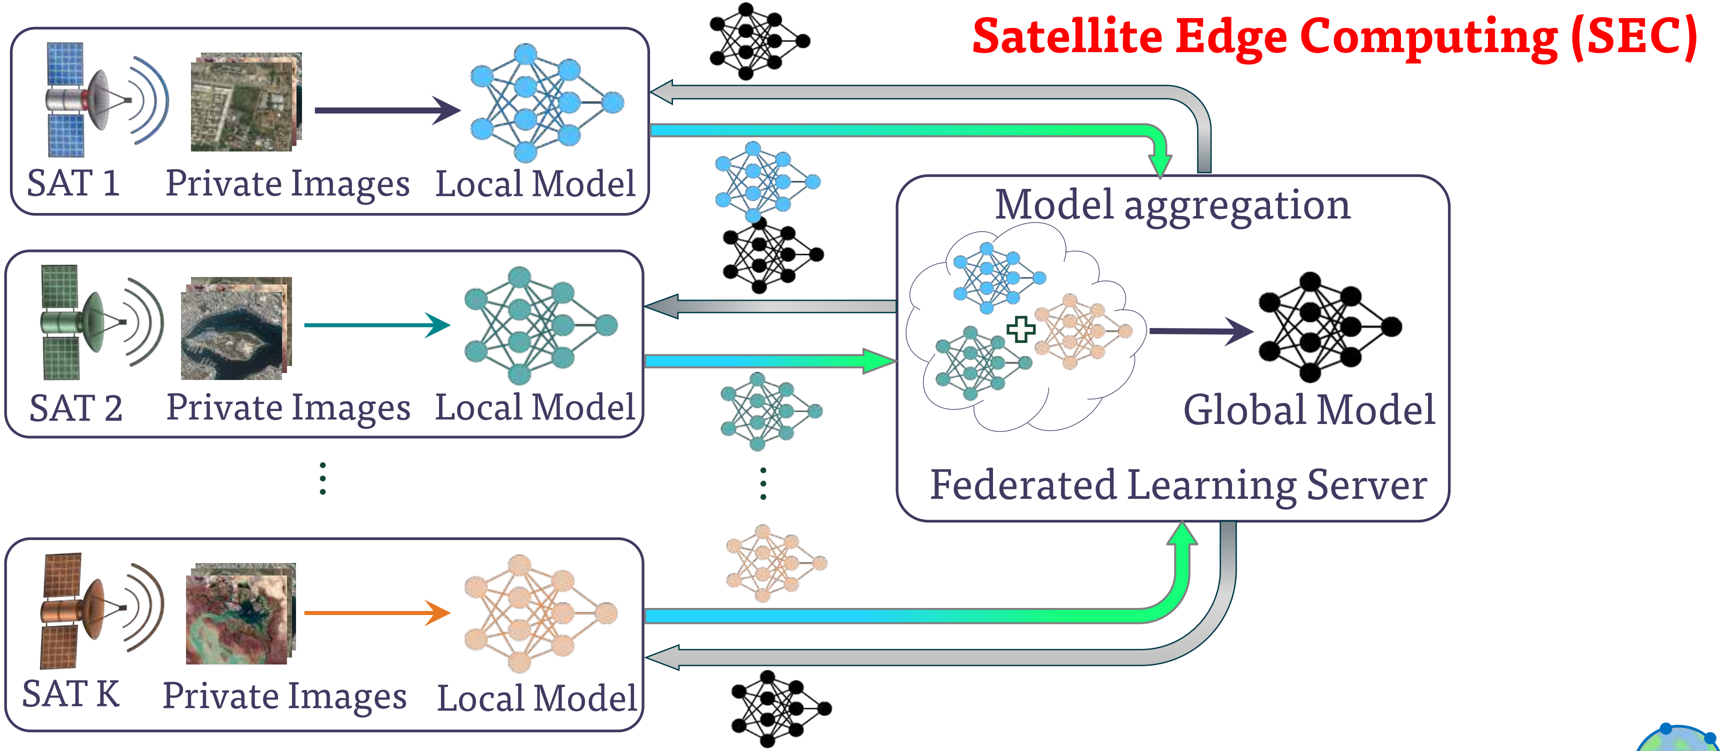

Challenges in SEC

- ❌ Limited computation and storage
  - ➢ LEO satellites cannot train large-scale ML models onboard
  - ➢ Overlooked by existing FL-LEO literature

- ❌ **Sporadic** and **irregular** visibility pattern
  - ➢ The iterative FL process will converge after several days or even weeks

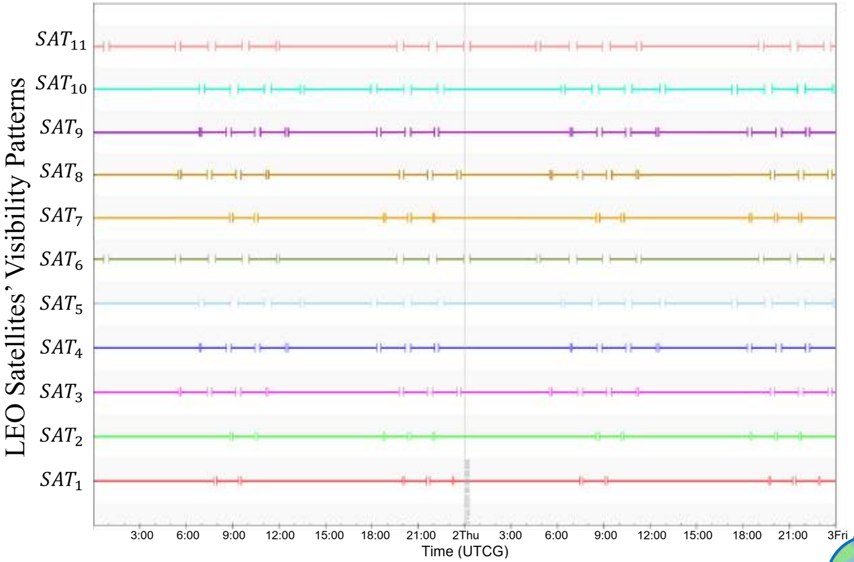

A visiting pattern over 2 days of a LEO network consisting of 11 satellites that communicate with a GS in Rolla, MO, USA

Contributions

- We propose an innovative framework FL to enable satellite edge computing (SEC)

Personalized Learning via Divide-and-Conquer

Enables satellites to train lightweight ML models, addressing the computation and memory limitations.

Contributions

- We propose an innovative framework FL to enable satellite edge computing (SEC)

Personalized Learning via Divide-and-Conquer → Orbital Model Retraining

→ Significantly reduces the number of communication rounds to accelerate convergence

Contributions

- We propose an innovative framework FL to enable satellite edge computing (SEC)

Personalized Learning via Divide-and-Conquer → Orbital Model Retraining → SEC Validation

We build a testbed that emulates SEC to validate our proposed solution

System Model

- LEO satellites collaboratively train a global ML model to minimize the global loss as:

Data size of satellite's n

$$\min_{w\in\mathbb{R}^d} F(w) = \sum_{n\in\mathcal{N}} \frac{\boxed{m_n}}{\boxed{m}} \boxed{F_n(w)}$$

Data size of all satellites

loss function of a satellite n

$$F_n(w) = \frac{1}{m_n} \sum_{x\in D_n} f_n(w; x)$$

- Each satellite n trains a local model by minimizing $F_n(\omega)$ through SGD:

Global iteration index

learning rate

$$w_n^{\boxed{\beta},j+1} = w_n^{\beta,j} - \frac{\boxed{\eta}}{\boxed{b}} \sum_{i=1}^{b} \nabla f_n(\boxed{w_n^{\beta,j}}; x_n^i), \quad i = 0, 1, 2, ..., I-1$$

Mini-batch size

Local model of satellite i at the j-the local iteration

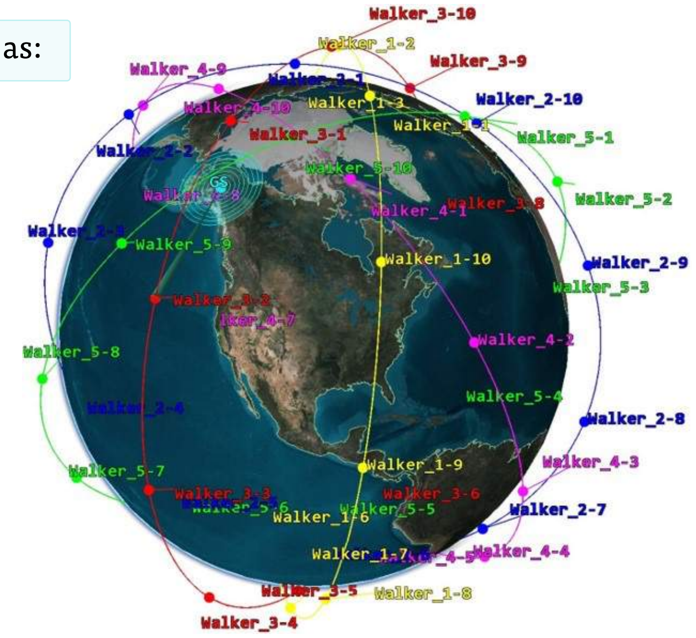

LEO network of N satellites

- The server aggregates all the satellites' local models into a global model:

$$\boxed{w^{\beta+1} = \sum_{n\in\mathcal{N}} \frac{m_n}{m} w_n^{\beta}}$$

1. Personalized Learning via Divide-and-Conquer

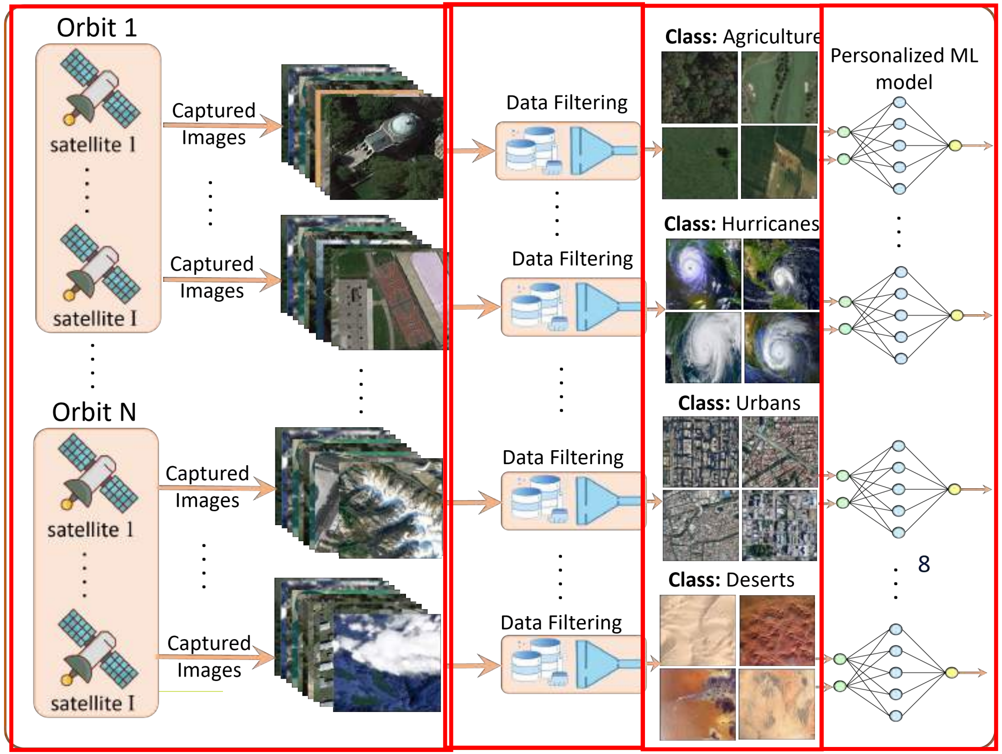

Converting multi-class classification problem into binary classification

- **Before Training**

✓ Each satellite filters its collected images for a single target class:

$$D_i^{\text{filtered}} = \{X_i \mid (X_i, y) \in D_i, \pi_i(y) = 1\}$$

✓ This converts the original complex multi-class problem into a binary classification problem.

$$t_{\text{train}}^i = \frac{\nu_i \cdot J \cdot c_{\text{process}} + c_{\text{overhead}}}{f_i}, \quad \text{where } \nu_i = \left\lceil \frac{m_i^{\text{filtered}}}{\kappa} \right\rceil$$

- **During training**

✓ Each satellite trains a personalized ML model by employing the one-vs-all strategy.

Orbital Model Retraining

2. Orbital Model Retraining

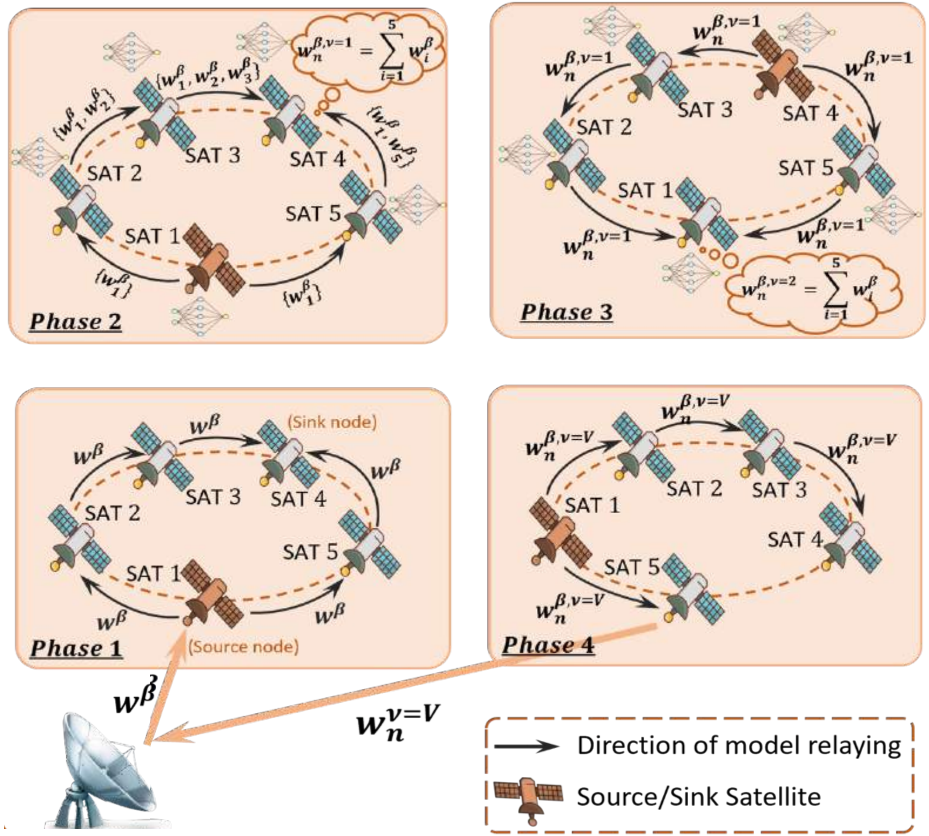

- **Objective**

Convert the binary classification tasks back to the original multi-class problem

- **Phase 1:**

Distribute global model to all satellites within each orbit

- **Phase 2**

Sink satellite aggregates all binary models received from satellites in the same orbit, forming an "orbital model"

- **Phase 3**

Retrain the orbital model for several orbital rounds: essentially FL in an extreme non-IID scenario

$$w_n^{\beta,v+1} = \sum_{i\in\mathcal{I}_n} \frac{m_i^{\text{filtered}}}{m_{\mathcal{I}_n}} w_i^{J,v}, \quad \text{where } m_{\mathcal{I}_n} = \sum_{i\in\mathcal{I}_n} m_i^{\text{filtered}}$$

- **Phase 4**

The first visible satellite to the FL server (ground station) transmits orbital model to server

Evaluation

Experimental Setup

- LEO Satellite Network
  - Two LEO networks of 60 satellites evenly distributed over 6 orbits, one with inclination of 45° and the other 5 with inclination of 85°
  - STK simulator was used to generate satellite visibility pattern
  - Ground station is positioned in Rolla, MO, USA as the FL server
- Satellite Training (Testbed)
  - Jetson Nano with NVIDIA Maxwell GPU to emulate each satellite
  - Dataset: EuroSat (high-resolution real-satellite images) containing 10 classes of land cover
  - Each satellite trains a VGG-16 model
  - Also uses MNIST, CIFAR-10, and CIFAR-100 datasets for comparison with SOTA

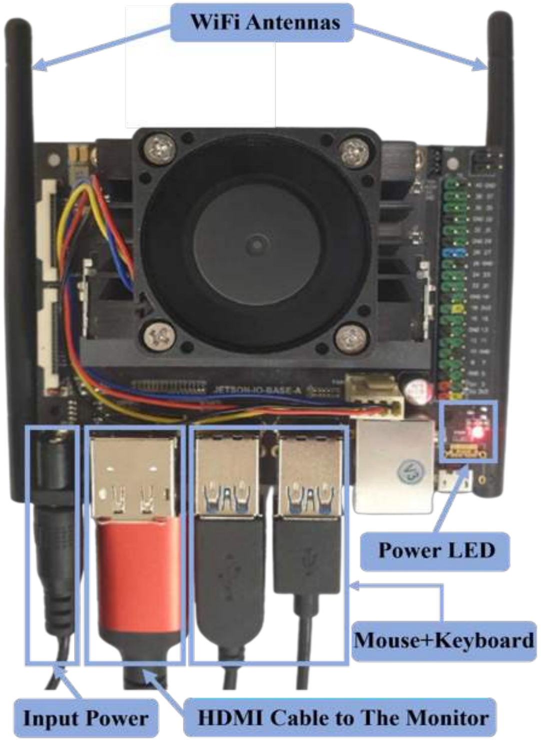

Evaluation

Results 1: Convergence

Comparing with SOTA

TABLE III: Comparison of convergence time and accuracy under non-IID settings (near-polar Constellation).

	FL-LEO Approaches	Accuracy (%)			Convergence time (h)
		MNIST	CIFAR-10	CIFAR-100	
Sync FL	FedAvg [3]	79.41	70.68	61.66	60 ($\mathcal{PS}$ located anywhere)
	FedISL [7]	82.76	73.62	66.57	8 ($\mathcal{PS}$ located at the NP)
	FedISL [7]	61.06	52.11	47.99	72 ($\mathcal{PS}$ located anywhere)
	NomaFedHAP [8]	82.73	77.36	62.81	24 ($\mathcal{PS}$ located anywhere)
Async FL	FedAsync [25]	70.36	61.81	56.37	48 ($\mathcal{PS}$ located anywhere)
	FedSpace [6]	52.67	39.41	36.04	72 (satellite uploads some of its data)
	AsyncFLEO [14]	79.49	69.88	61.43	9 (sink satellite has sufficient visible period)
	Ours	**94.64**	**89.69**	**82.65**	**2.13** ($\mathcal{PS}$ located anywhere)

Classification results (EuroSat)

Fig. 6: Twenty randomly selected images from a Eurostat test set of 5400 samples, illustrating the predicted vs. ground truth labels. Blue and red color represent correct and incorrect predictions, respectively.

Evaluation

Results 1: Convergence

Individual class accuracy for EuroSat

TABLE IV: Evaluation of our approach on EuroSat Dataset.

Class \ Metric	# of images	Near-polar constellation (85°) ACC(%)	PC (%)	RC (%)	F1 (%)	Inclined constellation (45°) ACC(%)	PC (%)	RC (%)	F1(%)
AnnualCrop	600	97.63	92.06	94.67	93.34	98.39	91.71	94.0	92.84
Forest	600	98.19	96.73	98.67	97.69	99.52	96.59	99.17	97.86
HerbaceousVegetation	600	99.32	93.05	93.67	93.36	98.5	92.12	94.5	93.33
Highway	500	99.61	97.74	95.0	96.35	99.31	97.93	94.60	96.24
Industrial	500	98.89	98.95	93.80	96.30	99.30	98.94	93.40	96.09
Pasture	400	99.23	96.42	94.25	95.32	99.33	96.67	94.25	95.44
PermanentCrop	500	97.46	95.14	90.0	92.50	98.54	94.50	89.40	91.88
Residential	600	99.01	91.45	99.83	95.46	99.04	92.15	99.83	95.84
River	500	99.56	98.17	96.60	97.38	99.44	97.37	96.60	96.99
SeaLake	600	99.11	98.66	98.0	98.33	99.63	99.15	97.50	98.32

Better visibility | Less visibility

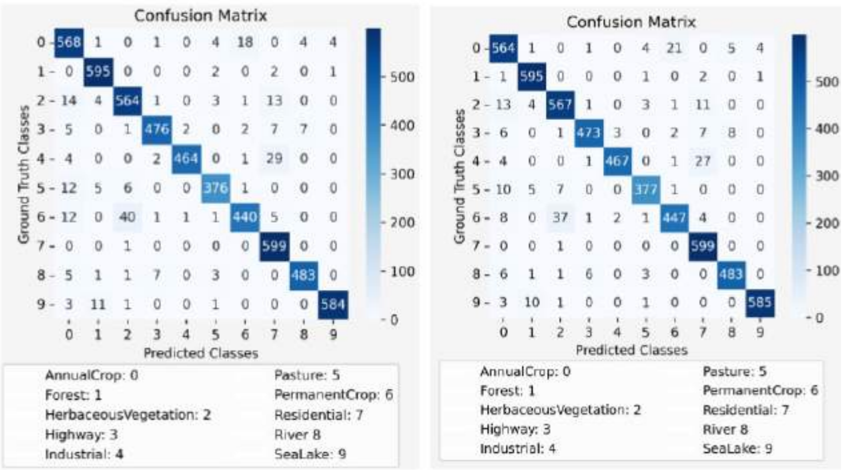

(a) Near-polar constellation (85°). (b) Inclined constellation (45°).

Fig. 5: Confusion matrix that compares 10 predicted and ground-truth classes for 5400 test images.

Evaluation

Binary Model Accuracy

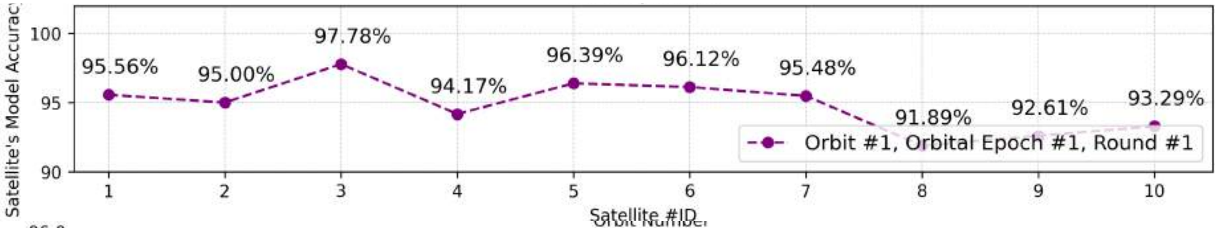

Orbital Model Accuracy

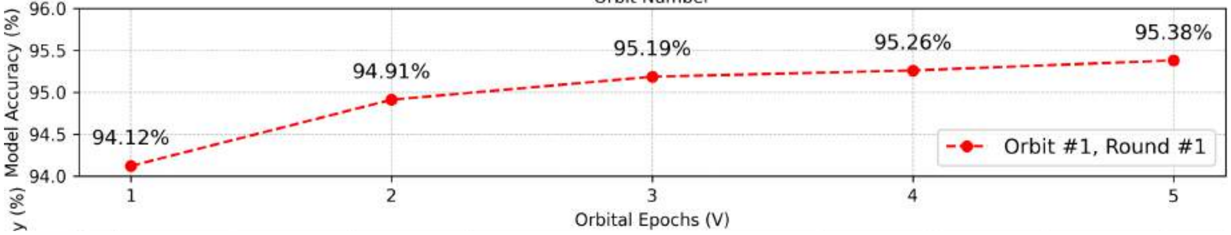

Global Model Accuracy

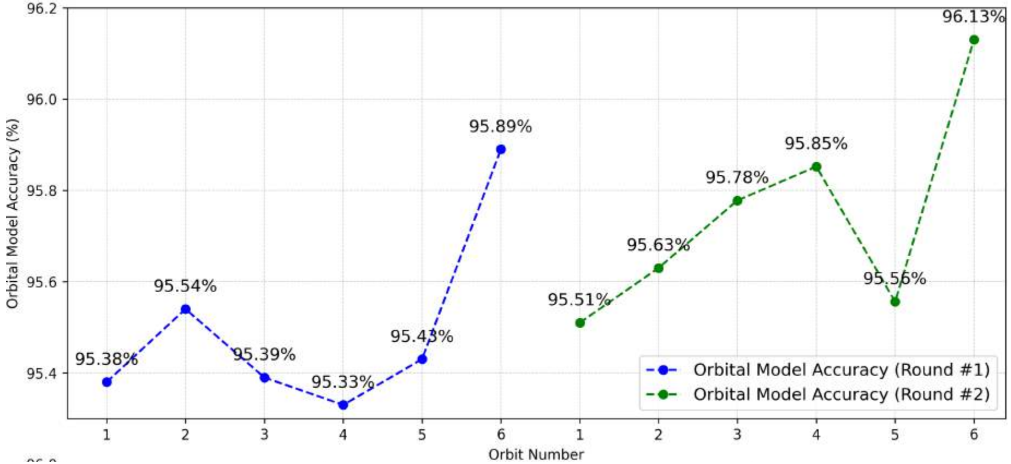

Results 2: Computation and Communication overhead; Energy Consumption

A. Computation Overhead

- Measured by floating-point operations per second (FLOPS)

B. Communication Overhead

- Measured by size of the global model

TABLE V: Comparison of computation and communication overheads of our approach in different settings.

(a) Computation overhead.

Model	FLOPS (G)
CNN using MNIST	11.91
CNN using CIFAR-10	15.58
CNN using CIFAR-100	28.13
VGG-16 using EuroSat	43.84

(b) Communication overhead.

Dataset	Size (MB)
MNIST	0.437
CIFAR-10	0.798
CIFAR-100	7.76
EuroSat	26.68

C. Energy Consumption

- Jetson Nano experienced GPU usage of 17--58% during training on various datasets
- Energy consumption ranged between merely 1.38--2.25 watts for each local model, demonstrating the advantage of being highly lightweight

From Lab to Space

- **Recent Launch** (March 6, 2024)
  - SpaceX launched an LEO satellite developed by Missouri University of Science and Technology, tasked with capturing Earth images.
  - Equipped with Raspberry Pi
- **Future Launch** (Feb. 2026)
  - Currently working on building two other LEO satellites: *Missouri Rolla* (MR) and *Missouri Rolla Second* (MRS).
  - MR and MRS will be equipped with Jetson Nano to train ML models onboard

Thank You

Questions?

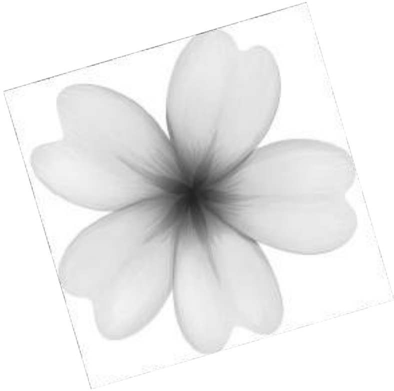

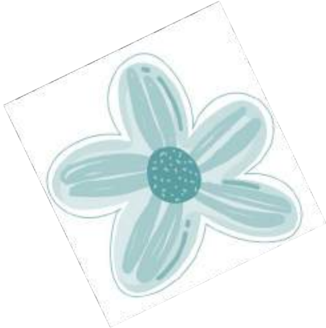

Take-Home Message

- AI-driven satellite edge computing (SEC) is coming.
- Federated learning integrated with LEO satellite networks (FL-LEO) is a promising approach but with challenges of computational and storage limitations
- We propose a solution to tackle this challenge in FL-LEO to enable SEC
  - Personalized ML model via divide-and-conquer
  - Orbital model retraining
- Validated by testbed

Funded by:

Acknowledgement:

Security & Privacy Preserving

All papers are available at https://tluocs.github.io

Computing & Communication

Convergence Speed & Efficiency