TECHNICAL UNIVERSITY OF CLUJ-NAPOCA

ACTA TECHNICA NAPOCENSIS

Series: Applied Mathematics, Mechanics, and Engineering
Vol. 69, Issue Special III, July, 2026

FEDERATED LEARNING FOR ON-ORBIT EARTH OBSERVATION ANALYTICS IN MILITARY SATELLITE CONSTELLATIONS

Simonel-Olimpiu DAVID, Andrei-Toma STOICA, Rareș-Cristian BIȘAG

Abstract: *Earth observation (EO) systems play a central role in modern military intelligence, surveillance, and reconnaissance by providing persistent, wide-area sensing of operational environments. However, conventional EO processing architectures rely heavily on ground-based analysis, leading to increased latency, high communication bandwidth requirements, and reduced resilience in contested or degraded communication scenarios. This paper explores federated learning as a system-level enabler for on-orbit EO analytics in distributed satellite constellations, allowing satellites to perform local inference and limited model adaptation while coordinating globally through the exchange of model updates rather than raw sensor data. We present a reference architecture for federated learning–enabled EO systems, analyze its operational advantages for time-critical military applications, and discuss key cybersecurity and trust challenges arising from distributed learning across space and ground segments. The paper argues that federated on-orbit analytics can significantly enhance responsiveness, robustness, and scalability of military EO systems when integrated within a secure, system-aware design framework.*

Keywords: *Earth Observation; Federated Learning; satellite constellations; spaceborne artificial intelligence; ISR systems; near real time EO analytics; contested communications; secure space systems*

1. INTRODUCTION

Military Earth observation (EO) systems are increasingly expected to support time-sensitive decision-making under conditions of uncertainty, limited connectivity, and adversarial interference. Applications such as early warning, surveillance of contested regions, monitoring of mobile threats, and rapid assessment of emerging events impose stringent requirements on latency, resilience, and autonomy. Traditional EO architectures, in which raw imagery is downlinked for centralized processing, struggle to meet these demands due to communication bottlenecks, ground station availability constraints, and the growing volume of data produced by modern sensor payloads.

In addition to volume and latency constraints, military EO systems must operate in an increasingly contested space environment, where communication links may be intermittently unavailable, deliberately degraded, or subject to denial. Reliance on continuous downlink and centralized processing introduces systemic vulnerabilities, as delays or disruptions at the communication or ground-segment level can directly affect the timeliness and reliability of intelligence products. These constraints motivate a shift toward greater autonomy at the space segment, where preliminary analysis and decision-support information can be generated even in the absence of immediate ground connectivity.

Recent advances in onboard computing and artificial intelligence have enabled satellites to perform inference directly in orbit, shifting parts of the perception pipeline closer to the sensor. However, standalone onboard models risk degradation over time due to changing environments, viewing geometries, and sensor conditions. Federated learning (FL) provides a

mechanism to address this limitation by enabling distributed model training across multiple platforms without requiring centralized access to raw data.

This paper examines the role of federated learning in enabling scalable, low-latency EO analytics within satellite constellations, proposing a system-level architecture and analyzing operational and cybersecurity implications using military EO systems as a representative operational context.

2. BACKGROUND AND RELATED WORK

Onboard artificial intelligence for EO has received growing attention as satellite platforms incorporate more capable processors and accelerators. Research and demonstrator missions have shown that tasks such as object detection, scene classification, and change detection can be executed directly on satellites, reducing downlink volume and enabling faster response cycles [1,2]. These efforts typically rely on pre-trained models uploaded from the ground and updated infrequently.

Federated learning, originally proposed as a communication-efficient training paradigm for distributed devices, allows multiple clients to collaboratively train a shared model by exchanging model updates rather than raw data [3]. While widely studied in terrestrial domains, federated learning in space-based systems introduces additional challenges related to intermittent connectivity, non-identically distributed data, heterogeneous hardware, and long platform lifetimes [4].

Applying federated learning in spaceborne systems differs fundamentally from terrestrial deployments. Satellite observations are strongly non-identically distributed due to orbital geometry, regional coverage, seasonal effects, and heterogeneous sensor payloads. Furthermore, participation in federated training rounds is inherently asynchronous, as satellites contribute updates only during available communication windows. Hardware constraints, radiation effects, and long mission lifetimes further limit the feasibility of frequent model retraining or complex optimization strategies. These characteristics necessitate federated learning designs that prioritize robustness, bounded adaptation, and predictable behavior over rapid convergence or maximal model expressiveness.

Prior work has explored federated learning for remote sensing and satellite networks primarily from an algorithmic standpoint, often assuming idealized communication conditions or homogeneous data distributions. In contrast, military EO systems operate under adversarial conditions, where communication links may be degraded or contested and where learning behavior must be tightly governed to ensure predictable system performance. This motivates a system-level analysis that integrates federated learning within the broader EO architecture rather than treating it as an isolated machine learning component.

3. SYSTEM ARCHITECTURE FOR FEDERATED ON-ORBIT EO ANALYTICS

The proposed architecture treats each satellite in the constellation as a federated learning client capable of performing local inference and limited model adaptation on recent EO observations. Optical or synthetic aperture radar data are preprocessed onboard to extract features or tiles relevant to the monitoring task, which are then used both for inference and for incremental local training.

Local model updates are periodically generated from recent observations and transmitted during available satellite communication windows. Rather than downlinking raw imagery, only encrypted and authenticated model updates are exchanged, reducing bandwidth requirements and limiting exposure of sensitive data. A ground-based federated aggregation service validates incoming updates, performs anomaly screening and robust aggregation, and maintains a versioned global model registry. Updated global models are redistributed to satellites during subsequent communication opportunities.

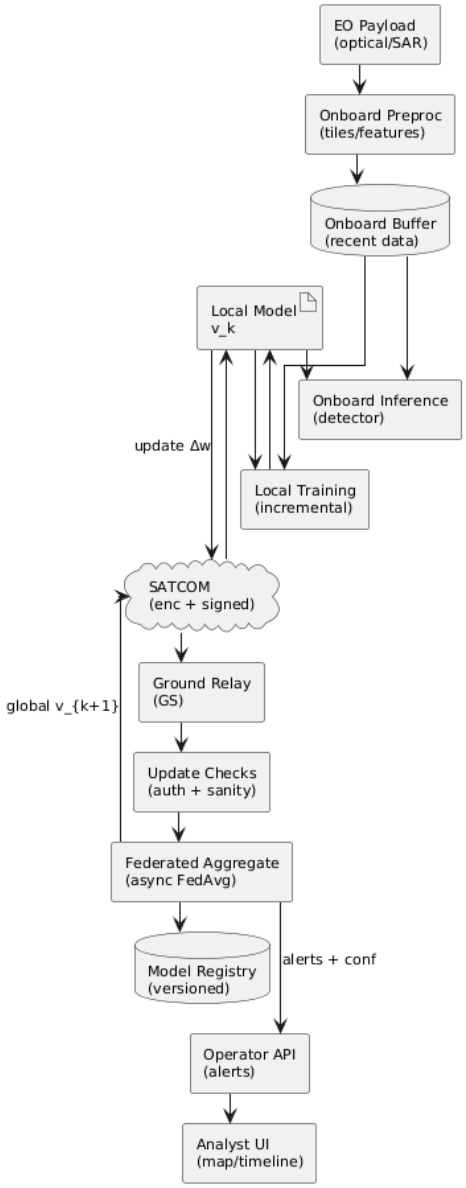


Fig. 1. Architecture for FL-enabled on-orbit EO analytics

Figure 1 Illustrates the end-to-end architecture, highlighting the vertical flow from onboard sensing and learning to ground-based aggregation and back to in-orbit inference. The architecture is intentionally modular, allowing integration with existing EO payloads and ground infrastructure while supporting asynchronous participation of satellites with intermittent connectivity.

Several design choices in the proposed architecture are driven by operational constraints specific to military EO systems. Local training is intentionally limited to incremental parameter updates rather than full retraining in order to bound model drift and ensure predictable inference behavior. Aggregation and validation are centralized on the ground segment, where greater computational resources and governance mechanisms can be applied before redistributing updated models. The architecture also assumes intermittent connectivity, avoiding synchronization barriers that would otherwise stall learning progress when subsets of the constellation are unavailable.

3.1 System Assumptions and Operational Constraints

The proposed architecture is designed under a set of explicit operational assumptions that reflect realistic military EO deployments. First, satellites are assumed to operate with limited onboard computational resources and strict power budgets, constraining the complexity of both inference and local training workloads. As a result, learning is restricted to lightweight, incremental updates rather than full retraining or large-batch optimization.

Second, communication opportunities between space and ground segments are assumed to be intermittent and asymmetric, with downlink capacity varying across orbital passes and operational theaters. The federated learning process is therefore designed to be asynchronous, tolerating partial participation and delayed updates without stalling system-wide operation.

Third, the architecture assumes that ground-based aggregation remains the primary locus for model governance, validation, and version control. While this introduces a degree of centralization, it enables the application of robust validation policies, auditing mechanisms, and controlled deployment strategies that are difficult to enforce in a fully decentralized space segment.

Finally, the system assumes that EO analytics provide decision-support information rather than authoritative targeting or engagement directives. This separation preserves alignment with established command

and control structures and ensures that on-orbit learning enhances situational awareness without embedding weapon-specific logic within the EO layer.

4. MILITARY USE CASES AND OPERATIONAL IMPLICATIONS

Federated on-orbit EO analytics provide several advantages for military applications where responsiveness and resilience are critical. By performing inference directly in orbit, satellites can generate preliminary alerts or localization cues with minimal latency, even when downlink opportunities are limited. These outputs can be ingested by downstream intelligence, surveillance, and reconnaissance (ISR) systems, air-defense infrastructures, or unmanned platforms for further verification and action.

A representative use case involves wide-area monitoring of regions with elevated risk of mobile or transient threats. On-orbit inference enables satellites to generate preliminary localization cues or change alerts shortly after acquisition, which can be forwarded to ground-based ISR systems or used to prioritize subsequent sensing actions. In such scenarios, federated learning allows the detection models to adapt gradually to region-specific backgrounds, infrastructure patterns, or environmental conditions, improving sensitivity while maintaining low false-alarm rates. Importantly, the EO system remains decoupled from downstream effectors, providing intelligence products rather than weapon-specific guidance.

The federated learning paradigm allows models to adapt gradually to regional conditions, sensor aging, and evolving operational environments without requiring centralized retraining on large datasets. This is particularly relevant in scenarios where data distributions vary significantly across theaters of operation or where rapid environmental changes would otherwise degrade model performance.

From an operational standpoint, federated learning also improves system robustness. Because learning and inference are distributed across the constellation, the system degrades gracefully under partial satellite loss or communication disruption. This aligns with military requirements for survivability and continuity of service in contested domains.

4.1 Integration with Multi-Domain ISR and Decision Cycles

In operational military contexts, EO systems rarely function in isolation, but rather as components of broader intelligence, surveillance, and reconnaissance (ISR) ecosystems spanning multiple domains. The outputs generated by federated on-orbit EO analytics—such as alerts, localization cues, or confidence-scored change indicators—are most valuable when integrated into existing decision cycles and fused with complementary data sources, including signals intelligence, airborne sensors, and ground-based observations.

By reducing the latency between acquisition and preliminary analysis, on-orbit inference enables EO systems to contribute earlier within the intelligence preparation and assessment process. This temporal advantage can support faster prioritization of sensing resources, adaptive tasking of other ISR assets, and more timely situational awareness updates, particularly in dynamic or rapidly evolving operational environments. Importantly, the federated learning architecture supports gradual model adaptation without disrupting these workflows, preserving consistency and interpretability of outputs across time and platforms.

From a systems perspective, the decoupling of EO analytics from downstream effectors remains a critical design principle. Federated on-orbit learning enhances the quality and timeliness of intelligence products, while leaving authority for decision-making, escalation, and engagement within established command and control structures. This separation supports interoperability, reduces integration risk, and aligns with doctrinal requirements for accountability and human oversight in military operations.

5. SECURITY AND TRUST CONSIDERATIONS

Introducing federated learning into satellite EO systems expands the attack surface by increasing the frequency and importance of communication between space and ground segments. While raw sensor data remain onboard, model updates themselves become sensitive assets that must be protected against interception, spoofing, replay, and manipulation.

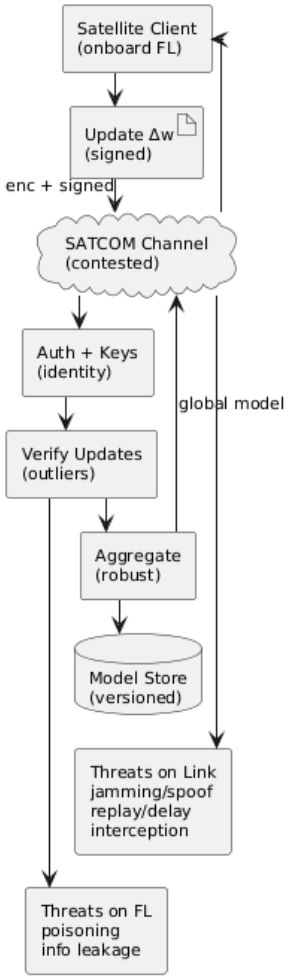


Fig. 2. Trust boundaries and threat model for FL in satellite-based EO systems

Adversarial attempts to inject malicious updates or selectively suppress legitimate ones can bias global models or degrade performance in specific regions. Additionally, inference attacks may extract information about local observation patterns from shared updates, even when communication is encrypted. These risks necessitate strong authentication mechanisms, cryptographic protection of update exchanges, and robust validation and anomaly detection of incoming contributions before aggregation.

Figure 2 presents a trust-boundary and threat model for federated satellite learning, identifying contested communication zones and the controls applied at each stage of the update pipeline. Rather than embedding weapon-specific logic within the EO system, the architecture focuses on delivering trustworthy, confidence-aware outputs that can be consumed by external military systems under established command and control frameworks.

Beyond individual attack vectors, federated learning introduces system-level governance challenges related to trust management across long-lived satellite platforms. Model update acceptance policies, participation thresholds, and rollback mechanisms must be defined to ensure that anomalous behavior does not propagate across the constellation. Versioning, audit trails, and controlled deployment of updated models are therefore essential components of any operational implementation. From a military perspective, these mechanisms support accountability, post-event analysis, and alignment with established command and control procedures.

6. CONCLUSION

Federated learning offers a promising path toward scalable, low-latency, and resilient Earth observation analytics in military satellite constellations. By enabling distributed learning and inference directly in orbit, EO systems can reduce dependence on continuous ground connectivity while maintaining adaptability to changing operational environments. However, realizing these benefits requires careful system-level design that integrates learning, communication, and security considerations from the outset. Future work should explore hybrid learning strategies, rigorous validation frameworks, and integration with multi-domain ISR architectures to further mature federated

on-orbit EO systems for operational deployment. While the proposed architecture highlights the potential of federated learning for on-orbit EO analytics, several limitations must be acknowledged. The scope of this work is intentionally focused on system-level design rather than algorithmic performance or quantitative evaluation, as the effectiveness of federated learning in spaceborne environments depends strongly on mission-specific factors such as sensor type, orbital configuration, and communication constraints. Moreover, bounded onboard training and conservative aggregation policies, while necessary for predictability and security, may limit the speed of model adaptation in highly dynamic environments. These trade-offs underscore the importance of tailoring federated learning strategies to specific operational contexts and of complementing on-orbit learning with periodic centralized retraining when communication conditions permit.

7. REFERENCES

[1] ESA Φ-lab. *Onboard Artificial Intelligence for Earth Observation Satellites*.

[2] Gorelick, N., et al. *Google Earth Engine: Planetary-scale geospatial analysis for everyone*, Remote Sensing of Environment, 2017.

[3] McMahan, B., et al. *Communication-Efficient Learning of Deep Networks from Decentralized Data*, AISTATS, 2017.

[4] Kairouz, P., et al. *Advances and Open Problems in Federated Learning*. Foundations and Trends in Machine Learning, 2021.

[5] Bonawitz, K., et al. *Towards Federated Learning at Scale: System Design*. SysML, 2019.

[6] Bhagoji, A. N., et al. *Model Poisoning Attacks in Federated Learning*. Secure and Trustworthy ML Workshop, 2019.

[7] ECSS. *Space System Security Engineering: Cybersecurity for Space Missions*.

Învățare federativă pentru analize de observație a Pământului realizate la bordul sateliților în constelații militare

Rezumat*:* Observația Pământului (EO) joacă un rol esențial în sistemele moderne de informații, supraveghere și recunoaștere militare, oferind capacități de senzare persistentă și la scară largă asupra mediilor operaționale. Cu toate acestea, arhitecturile convenționale de procesare EO se bazează în mare măsură pe analiza realizată la sol, ceea ce conduce la latențe ridicate, cerințe mari de lățime de bandă pentru comunicații și o reziliență redusă în scenarii cu comunicații degradate sau contestate. Această lucrare analizează utilizarea învățării federative ca mecanism de nivel sistem pentru realizarea analizelor EO la bordul sateliților, în cadrul constelațiilor distribuite, permițând efectuarea inferenței locale și a unei adaptări limitate a modelelor direct în orbită, concomitent cu coordonarea globală prin schimbul de actualizări de model în locul datelor brute de senzare. Este prezentată o arhitectură de referință pentru sisteme EO bazate pe învățare federativă, sunt analizate avantajele operaționale pentru aplicații militare sensibile la timp și sunt discutate principalele provocări legate de securitatea cibernetică și de managementul încrederii rezultate din învățarea distribuită între segmentele spațiale și terestre. Lucrarea susține că analizele EO realizate on-orbit, integrate într-un cadru de proiectare sigur și orientat pe sistem, pot îmbunătăți semnificativ capacitatea de reacție, robustețea și scalabilitatea sistemelor EO militare.

Simonel-Olimpiu DAVID, MSc Student, Babeș-Bolyai University, Faculty of Mathematics and Computer Science, Distributed Systems in Internet Masters, simonel.david@stud.ubbcluj.ro., +40 264 405 300, 1 Mihail Kogalniceanu Street, Cluj-Napoca, Romania.

Andrei-Toma STOICA, PhD candidate, National University of Science and Technology POLITEHNICA Bucharest, Faculty of Automatic Control and Computers, Computers and Information Technology, andrei_toma.stoica@stud.acs.upb.ro, +40 214 029 100, 313 Splaiul Independentei, Bucharest, Romania.

Rareș-Cristian BIȘAG, PhD Student Industrial Engineering, National University of Science and Technology POLITEHNICA Bucharest (UNSTPB), Faculty of Entrepreneurship, Business Engineering and Management, rares.bisag@upb.ro, +40 214 029 100, 313 Splaiul Independentei, Bucharest, Romania.