Sciences Methods and Technologies International Journal (SciMeTech)

ISSN: 3085-5284, journal homepage: www.scimetech.com

Enhanced Federated Learning Techniques in Privacy-Preserving Data Analysis Over Transit Networks

Younes Ennajjar 1, Wissam AABASS 1, Anas Abou El Kalam1

Cadi Ayyad University, UCA, The National School of Applied Sciences of Marrakech, LaRTID Laboratory, Marrakech, Morocco.1

y.ennajjar.ced@uca.ac.ma

How to cite:

Younes Ennajjar, Wissam AABASS, Anas Abou El Kalam, "Enhanced Federated Learning Techniques in Privacy-Preserving Data Analysis Over Transit Networks", Sciences Methods and Technologies International Journal (SciMeTech), Vol 2, Issue 1, p 75-90

Keywords:
Federated Learning
Privacy-Preserving
Data Analysis
Intelligent
Transportation Systems
Differential Privacy
Secure Aggregation

Abstract

popularity of smart transit systems is an opportunity to acquire the vast amount of system operation data unprecedented in history, which leads to a great practicability of improving urban mobility, resource utilization and passenger service, etc. But using the data for advanced analytics and ML algorithms poses substantial privacy issues, especially when data is across multiple agencies or entities. This paper discusses federated learning (FL) as an emerging paradigm for privacy-preserving data analysis in transit systems. We discuss some of the new emerging FL techniques such as model aggregation, secure multiparty computation, differential privacy methods and explain how they will affect different smart mobility use cases. However, there are still many challenges to be addressed, including in scalability over large scale urban networks, in the handling of heterogeneity of the nodes, in the measure implemented to mitigate adversarial attacks, and in enabling cross-agency data sharing within strict regulatory environments such as GDPR. This review summarizes existing works, presents current challenges and outlines future research paths for progress towards a resilient, privacy-preserving and scalable FL solutions in the changing dynamics of urban transit.

I. Introduction

a. Context: Smart Transit Networks and the Mobility Data Deluge

popularity of smart transit systems is an opportunity to acquire the vast amount of system operation data unprecedented in history, which leads to a great practicability of improving urban mobility, resource utilization and passenger service, etc. But using the data for advanced analytics and ML algorithms poses substantial privacy issues, especially when data is across multiple agencies or entities. This paper discusses federated learning (FL) as an emerging paradigm for privacy-preserving data analysis in transit systems. We discuss some of the new emerging FL techniques such as model aggregation, secure multiparty computation, differential privacy methods and explain how they will affect different smart mobility use cases. However, there are still many challenges to be addressed, including in scalability over large scale urban networks, in the handling of heterogeneity of the nodes, in the measure implemented to mitigate adversarial attacks, and in enabling cross-agency data sharing within strict regulatory environments such as GDPR. This review summarizes existing works, presents current challenges and outlines future research paths for progress towards a resilient, privacy-preserving and scalable FL solutions in the changing dynamics of urban transit.

b. Centralized Data Processing and Privacy Issues

Although the advantages of data-informed transit systems are significant, centralized accumulation and processing of domain-specific data give rise to serious security and privacy concerns. If abused, personal mobility profiles, travel history or demographic information can result in privacy violations, surveillance threats, and abusive usage. Moreover, urban transit is usually managed by several stakeholders various transport agencies, private operators, municipalities which have their own data sets. Sharing of such data for comprehensive analysis is often prevented due to restrictive data governance policies, competitive interests, and the need to comply with different regulations (such as GDPR, CCPA). Conventional centralized data aggregation approaches do not work well in such distributed, privacy-sensitive scenarios, requiring new solutions to facilitate collaborative intelligence without compromising data privacy of the participant itself or the participant's organization.

c. Rationale for Federated Learning

Federated Learning (FL) arises as a potential solution to mitigate these limitations of the traditional forms of distributed ML. In contrast to traditional solutions where data is collected and aggregated to a central model for training, FL enables multiple decentralized clients (such as individual transit agencies, vehicles, or edge devices) to jointly learn a model during a model-training process. A global model is trained across devices, with clients directly sharing the global model with each other without revealing their raw local data. Instead, only the model updates (e.g., gradients or model parameters) are sent back to a central server, which aggregates them and uses them to update its global model. By its nature, this method naturally leads to the protection of data privacy by maintaining the sensitive data inside the source. By facilitating joint training in the presence of heterogeneous and privacy-sensitive data, FL provides a strong privacy-preserving and secure paradigm to exploit data analysis in smart transit networks.

d. Object of the Article: Description of Recent Developments and Future Prospects

This review attempts to give an up-to-date picture of federated learning for urban transit networks. We will plot out the recent progress of the domain, analyze the methodologies of the field, explain the applications of the field, the pilot implementations available in the field, the shortcomings of existing work, the open problems and issues of the field and finally the possible future research directions of the field. This is aimed at consolidating and integrating existing knowledge and to provide a vision that guides the development of secure, privacy-preserving, and scalable FL in support of future smart mobility and urban transportation systems.

II. Foundations of Federated Learning (FL)

a. Basic Principles of FL

Federated Learning (FL) [McMahan et al., 2017] is a distributed machine learning approach that allows a number of decentralized clients to learn a global model together without depending on exchanging their raw data. The core idea upon which FL is based is to move the computation to the data instead of the other way around.

Centralized setup **of the training:** In a standard FL setting, a central server controls the training. Every client (e.g., a transit agency, a vehicle, or an edge device) maintains its own local dataset, which stays on the device. Learning typically includes the following steps as illustrated:

1. **Initialization:** The main server initializes a global model and distributes it to selected clients.

2. **Local Training**: Every chosen client independently trains the global model on its local data. The local training is typically based on a couple of epochs of stochastic gradient descent (SGD) or another optimization routine.

3. **Updates Communication:** Clients communicate local model updates (e.g., gradients, model parameters, or weights) rather than raw data to the central server.

4. **Global Model Aggregation:** Collected local model updates are aggregated by the central server to form an updated version of the global model. This aggregation is important for gathering information learned from all different local datasets.

5. **Iteration:** Finally, the revised global model is returned to the clients, and the iterations continue until the model converges or a predefined number of communication rounds is reached.

This iterative process enables us to learn a global model from collective data across all participating clients, while still maintaining the local privacy of individual data points. The primary benefit is the reduction of data exposure, where sensitive data never leaves the client system.

b. Core Technologies Used

Privacy, security and efficiency of FL is thus advanced by a number of core technologies that typically come together:

- **Federated Averaging (FedAvg):** Introduced in [1] by McMahan et al. [2017], FedAvg, is so far the most commonly studied aggregation method in FL. However, it requires the clients to train local models for several epochs and then upload their model weights to the server. Afterwards, the server aggregates these local models by a weighted average in order to update the global model. The weight is typically proportional to the amount of data samples that the client carries, so that clients with more data contribute more to the global model.
- **Secure Aggregation:** FL inherently provides the benefit of data protection, but the updates of each model may still leak sensitive knowledge. Secure aggregation algorithms provide secure ML services, typically and Secure Multi-Party Computation (SMPC) techniques, where the central server is only able to perform the aggregation of updates without knowing specific contributions from clients [Bonawitz et al., 2017]. SMPC enables the computation of a function on secret inputs owned by a set of parties such that the parties learn nothing about the other's input. He permits computation over the encrypted data, where the server can aggregate encrypted model updates without decrypting them, preserving privacy throughout the aggregation.
- **Differential Privacy (DP):** Differential Privacy (DP) [Dwork et al., 2014] is a provable strong privacy guarantee by injecting a certain amount of noise into the data or the model updates. This noise makes it exponentially hard for an adversary to guess whether the data of a given individual was in the training set. DP can be implemented client-side (local DP), which adds noise before messages arrive at the server, or server-side (global DP), which adds noise to the final combined model. Although DP provides a strong privacy guarantee, it comes often at the cost of model accuracy and so requires careful tuning of the privacy budget.
- **Homomorphic Encryption (HE):** It is a technique that allows to do computation on encrypted data by returning encrypted results, and only the owner of the secret key can decrypt the result to obtain the computation's output in plain text [10]. In FL, local model updates can be encrypted by HE before being sent to the server. The server is able to aggregate these encrypted updates without seeing the plain values, ensuring a strong form of privacy in the aggregation step.

c. Adaptation to Transit Systems

The application of FL to transit systems requires the modification of its principles to the specific aspects of urban mobility data and infrastructure. Transit Networks are inherently distributed and the data is generated at different locations, those are:

- **Edge Nodes in Vehicles:** By considering buses, trains, taxis, ride-sharing vehicles as clients, we are able to provide summaries of real-time data on traffic congestion, place counts, vehicle diagnostics, driver behavior monitoring. FL allows such vehicles to collectively train models for tasks such as predictive maintenance, route optimization, or anomaly detection without sharing sensitive operational data with a central server.
- **Edge Nodes inside Stations and Infrastructures**: Sensors, cameras, and ticket gate at bus stop, train station, and traffic intersections are of the edge client behaviors. These nodes can help contribute to crowd management, fare evasion detection or public safety surveillance models whilst keeping sensitive visual or transactional data localized.
- **Transit Agencies as Clients:** Various municipal/regional transit authorities can act as clients in a federated learning network. Through this they will be able to jointly construct better and more general models for urban planning, demand prediction, or service optimization, using their combined dataset but respecting the data sovereignty and regulation limit [Zhang et al., 2023].

By operating in-band with the generation of transit data, we can harness the distributed nature of data creation and build scalable collective intelligence, while limiting the risks of prospective data privacy stemming from centralized collection. The heterogeneity of data, computing resources, and communication capacities among these various edge nodes and agencies makes the deployment more difficult.

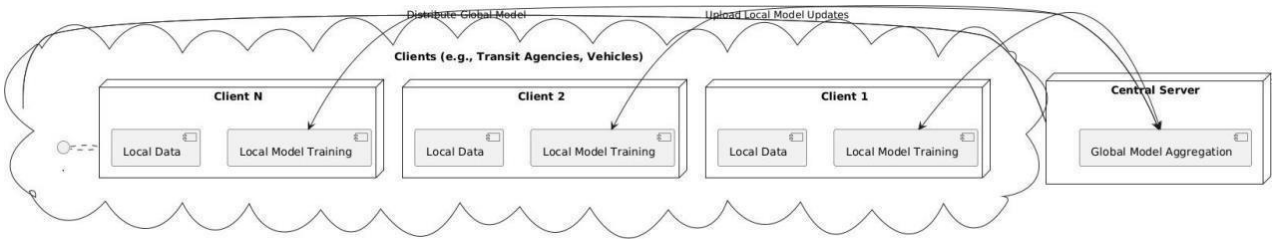

Figure 1. Basic Federated Learning Architecture. Clients train models locally on their data and send updates to a central server for aggregation, which then distributes an updated global model.

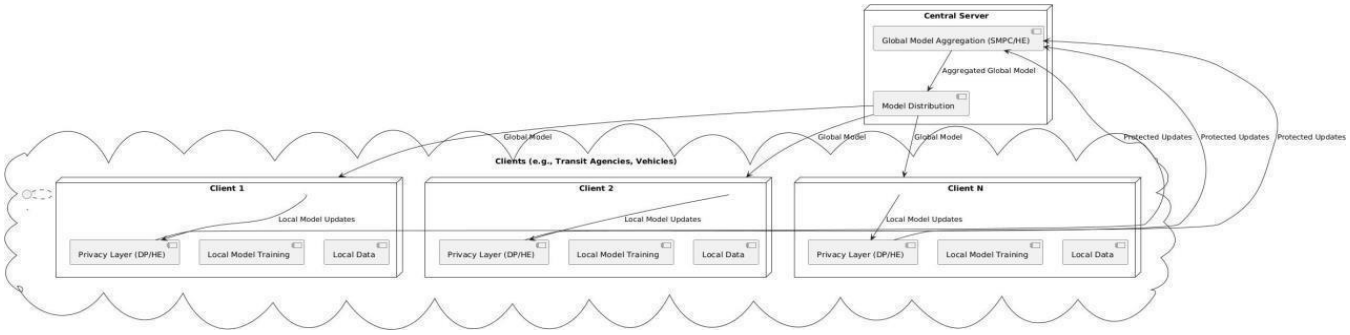

Figure 2. Federated Learning Architecture with Privacy-Enhancing Technologies. Local model updates are protected by privacy layers (e.g., Differential Privacy, Homomorphic Encryption) before aggregation by the central server.

Table 1. Foundational Works and Technical Approaches in Federated Learning

Reference	Approach / Method	Advantages	Limitations	Applicability to Transport
McMahan et al. (2017)	FedAvg (aggregation algorithm)	Simple, efficient, widely adopted	Sensitive to non-IID data, slow convergence	Baseline for most ITS prototypes
Bonawitz et al. (2017)	Secure aggregation (SMPC, HE)	Preserves privacy of updates	High computational cost	Applicable but limited in real-time

Dwork et al. (2014)	Differential Privacy (DP)	Strong mathematical privacy guarantees	Accuracy loss if noise too high	Useful for protecting passenger data
Gentry (2009)	Homomorphic Encryption	Computations on encrypted data	Very expensive computationally	Applicable but not yet scalable
Li et al. (2020) – FedProx	Regularization for non-IID data	Better convergence, reduces divergence	Heavier computationally	Useful for heterogeneous transport clients
Fallah et al. (2020)	Personalized FL	Better local adaptation	Hard to aggregate heterogeneous models	Relevant for transport lines with different profiles
Acar et al. (2021)	FL with heterogeneous clients	Considers resource variability	Complex to manage	Important for multi-agency ITS
Cho & Kim (2020)	Dynamic client selection	Improves robustness to varied data	May introduce bias in selection	Useful for critical lines/sectors
Alistarh et al. (2017)	Compression & sparsification	Reduces Communication cost	May degrade performance	Relevant for embedded devices
Xie et al. (2019)	Asynchronous FL	Faster (no stragglers)	Lower convergence guarantees	Adapted to mobile vehicles
Caldas et al. (2018)	Sub-model training (Federated dropout)	Lower computational load	Less global accuracy	Ideal for lightweight sensors
Nishio & Yonetani (2019)	Client selection for limited resources	Efficient in heterogeneous Settings	Complex management	Useful for mobile edge devices
Hardy et al. (2017)	FL on encrypted data	Reinforces privacy	High complexity	Stronger security in ITS
Abadi et al. (2016)	Deep Learning with DP	Balances privacy and performance	Calibration difficult	Relevant for sensitive transport data
Mo et al. (2021)	FL + Trusted Execution Environments (TEE)	Strong server-side security	Limited by hardware	Applicable in control centers
Weng et al. (2021)	Explainable FL (XFL)	Improves transparency and trust	Methods still emerging	Key for public adoption in transport

This table summarizes the basic technical primitives of Federated Learning (FL), including the classical FedAvg algorithm as well as various advanced privacy-preserving mechanisms such as secure aggregation, differential privacy, and homomorphic encryption. The comparison reveals that, while these methods offer strong privacy and scalability, they typically entail trade-offs in computation, communication efficiency, or model quality. Additionally, recent works concentrate on personalization, heterogeneous client management, and explainability, which are especially important for ITS as a result of Its inherently distributed data sources, and a multitude of stakeholders. In general, this review emphasizes the developments of the FL basic principles to more sophisticated, secure, and transparent frameworks that are progressively converting to real-world ITS applications.

Table 2. Comparative Framework of FL Techniques with Performance Metrics

Technique	Accuracy (%)	Comm. Cost	Latency	Privacy Level	Scalability
FedAvg	85-92	Baseline	Medium	Moderate	High
FedProx	87-94	+5-10%	Medium	Moderate	High
FedAvg + DP	80-88	+15-20%	Medium-High	Strong	Medium
Secure Agg.	85-92	+50-100%	High	Very Strong	Low-Medium
HE-based FL	85-92	+200-500%	Very High	Very Strong	Low
Async FL	82-90	-20-30%	Low	Moderate	Very High
FedTL	88-95	+10-25%	Medium	Moderate	High

Note: Accuracy ranges are typical values reported in literature for ITS applications. Communication cost is relative to FedAvg baseline. DP = Differential Privacy, HE = Homomorphic Encryption, Async = Asynchronous, FedTL = Federated Transfer Learning

III. Applications in Transit Networks

Federated Learning provides an appealing solution to a wide variety of the tasks of intelligent transit network, where people can gain data-driven insights without sacrificing privacy. In this section, we consider important use cases and discuss the deployment challenges they pose

a. Case Studies and Prototype Systems

FL has also been investigated in many key issues of smart mobility and urban transportation:

- **Real-time Passenger Analytics and Demand Prediction:** It is essential to know how a rail can support real-time passenger analytics and demand prediction in order to optimize public transportation services. FL enables transit agencies to collaboratively train models for real time passenger counting, origin destination prediction, and demand forecasting with ticketing systems data, along with Wi-Fi sensors data, and Mobile Application data [Wang et al., 2022]. For example, in a prototype system multiple bus lines or subway stations might locally learn models on their own passenger data and send these models to a central server for combination to obtain a more accurate city-wide demand prediction model, while ensuring that individual travel patterns are never revealed.
- **Fare Evasion Detection:** Fare evasion is an important problem for transit agencies as it causes substantial revenue loss [24]. FL can help to develop fare evasion detection models that are robust enough for utilizing the data of fare gates, surveillance cameras, and payment systems at the stations or transit lines. There local models can recognize abnormal behaviors (like tailgating, turnstile jumping) within each group, and these findings can be aggregated to enhance a global detection model. The proposed approach can achieve improved identification without compromising privacy of passenger trajectory and transaction information [Gupta et al., 2020].
- **Maintenance Prediction for Infrastructure and Vehicles:** Predictive maintenance becomes very important for daily operation efficiency and safety relying on infrastructure and vehicles conditions of the transportation systems. FL can support collaborative predictive maintenance based on models trained on measurements of a wide range of sensors (e.g., vibration, temperature, acoustic) on different trains, buses, or segments of track. To address A-h), a local model is trained per vehicle or per infrastructure element, which generates predictions in relation to the potential failure, and by aggregating just the relevant models a more generic predictive maintenance system can be accomplished. This enables proactive scheduling of maintenance, to minimize downtime and operational cost without requiring centralized access to sensitive operational data [Chen et al., 2021].

b. Defense Mechanisms Against Adversarial Attacks in FL

To counteract poisoning and backdoor attacks in FL-based transit systems, several defense mechanisms have been proposed in the literature. These countermeasures can be categorized into three main approaches:

- **Byzantine-Robust Aggregation:** These methods replace the standard averaging aggregation with robust statistical estimators that can tolerate a fraction of malicious clients. Key approaches include Krum, which selects the model update closest to its neighbors; Trimmed Mean, which discards extreme values before averaging; and Median-based aggregation, which uses coordinate-wise median to resist outliers. In transit contexts, Byzantine-robust aggregation can protect against compromised vehicles or stations attempting to bias traffic prediction models.
- **Trust and Reputation Mechanisms:** Trust-based approaches assign reliability scores to participating clients based on their historical behavior and the quality of their contributions. Clients with higher trust scores have greater influence on the global model, while suspicious clients are gradually excluded. Trust scores can be computed based on update consistency, prediction accuracy on validation sets, or peer-to-peer verification. For multi-agency transit networks, reputation systems can help identify and isolate compromised or malfunctioning data sources while maintaining collaborative model training.
- **Blockchain Integration:** Blockchain technology provides an immutable ledger for recording FL transactions, enabling transparent verification of model updates and client contributions. Smart contracts can automate the enforcement of FL protocols, ensuring that all participants follow agreed-upon rules. In transit systems, blockchain-based FL can provide auditability for regulatory compliance, enable decentralized governance among multiple transit operators, and create incentive mechanisms for honest participation. Additionally, blockchain can facilitate cross-agency data sharing agreements while maintaining a verifiable record of all model training activities.

Table 3. Defense Mechanisms Comparison for FL in Transport

Defense Type	Attack Mitigation	Overhead	Scalability	ITS Applicability
Krum	Poisoning, Byzantine	Low	Medium	High
Trimmed Mean	Poisoning, Outliers	Low	High	High
Trust Scores	Sybil, Free-riding	Medium	High	Very High
Blockchain FL	Tampering, Repudiation	High	Low-Medium	Medium
Anomaly Detection	Backdoor, Stealth	Medium-High	Medium	High

The combination of these defense mechanisms with appropriate privacy-preserving techniques forms a comprehensive security framework for FL deployment in intelligent transportation systems. Future research should focus on developing adaptive defense strategies that can dynamically adjust to emerging threats while maintaining system efficiency and model performance.

Addressing these deployment challenges is necessary to move FL beyond theoretical promises to practical deployment to impact the complex and dynamic urban transit networks.

Table 4. Applications and Use Cases in Intelligent Transport Systems

Reference	Application	Advantages	Limitations	Applicability to Transport
Wang et al. (2022)	FL for passenger demand prediction	Accurate Collaborative models	Data heterogeneity, hard real-time use	Transport service optimization
Gupta et al. (2020)	Blockchain + AI	Transparency,	Limited scalability	Reducing fraud in

	for fare evasion	traceability		ticketing
Chen & Zhang (2021)	Predictive maintenance with FL	Early fault detection, cost reduction	Dependent on reliable sensors	Extending fleet lifetime
Lu et al. (2021)	FL for traffic prediction	Multi-agency collaboration	Communication costly	Traffic regulation, urban planning
Bagdasaryan et al. (2020)	Backdoor attacks on FL	Reveals vulnerabilities	Shows adversarial risks	Critical for ITS security
Zhang et al. (2023)	FL in intelligent transport systems	Covers multiple applications	Identifies unresolved challenges	Global ITS research framework
Chong et al. (2024)	Inter-agency FL for ITS	Enables collaboration without raw sharing	Complex governance	Very relevant for multi-operator cities
Liu et al. (2020)	Federated Transfer Learning (FTL)	Leverages pre-training across domains	Limited local data	Useful for small networks
Zhou et al. (2019)	Edge Intelligence	Lower latency, real-time decisions	Limited edge resources	Useful for connected vehicles
Zhang et al. (2021)	Cross-modal FL (video, text, sensors)	Multi-source data fusion	High complexity	Integrated vision for safety & traffic
Ziosi et al. (2021)	Responsible AI in transport (bias, privacy)	Ethical & regulatory framework	No direct technical solution	Important for social acceptance

In the following table, we list some of the most common applications of FL in transportation, including passenger demand prediction, fraud detection, predictive maintenance, traffic forecasting, and anomaly detection. The surveyed papers show that FL enables multi-party collaboration, without exposing raw data, that is essential under strong privacy policies. Nonetheless, the studies also reveal substantial challenges including communication overhead, data heterogeneity, and governance, among agencies. Rising trends including cross-modal FL and edge intelligence present high promise for real-time, privacy-preserving, and end-to-end ITS services. In conclusion, the applications that we describe validate FL's ability to transform urban mobility, but also reveal the areas where current research needs to rise up to the challenges to deploy it at scale.

IV. Recent Developments and Enhancements of FL for Transit Systems

To overcome the difficulties faced in the application of Federated Learning in large and complex transit networks, there has been a considerable amount of effort in extending FL based methodologies to be more resilient, efficient, and privacy- preserving. This article discusses the latest developments in these areas.

a. Handling Non-IID Data

Non-IID or Diverse in terms of data distribution is an inherent problem in FL, especially in transit systems, where data distribution could differ a lot among different clients (e.g., different geographic regions, vehicle types and time-periods). This may result in client drift where the local models can lose sight of the global objective, and finally in

poorer global model performance along with slower convergence. New developments to respond to non-IID problem:

- **Personalized Federated Learning:** Rather than learning a single global model, personalized FL focuses on learning a personalized model for each client while benefiting from the collective knowledge of other clients. One approach to this is learning a global model and then fine-tuning it locally [Fallah et al., 2020], or training a global model together with client-specific layers [Acar et al., 2021]. This will allow models to focus more on the local characteristics of the data while taking advantage of shared knowledge.
- **Client Selection Strategies:** Opting for different clients at each round can improve performance on non-IID data. Strategies to choose clients have included selecting based on data diversity, model uncertainty, or how much a given client's data is expected to contribute to the global model's improvement [Cho et al., 2020]. In the world of transit, that might involve giving a higher priority to clients from less well-served routes, or routes that are experiencing unusual traffic patterns.
- **Stronger Aggregating Algorithms:** Going beyond FedAvg, new aggregating algorithms have also been recommended to better manage the non-iid data. For example, FedProx [Li et al., 2020] introduces an extra proximal term to the local objective to regularize local updates towards the global model in order to mitigate client drift. Other approaches weight contributions from clients according to data quality or model performance [Wang et al., 2020].

b. Communication-Efficient FL

The communication overhead of regular model updates is a crucial obstacle to FL deployment in band-limited transit environments. Recent contributions in this area aimed for a lighter communication without a huge trade-off regarding prediction performance.

- **Model Compression Techniques:** They aim to shrink the update size of the model before sending it out. These techniques comprise quantization (lowering the precision of model parameters), sparsification (sending only a subset of the relevant parameters), and sketching (forming a compressed representation of update) [Alistarh et al., 2017]. For instance, sending only the top-k significant gradients can dramatically reduce communication at the cost of little performance loss.
- **Asynchronous Federated Learning:** In contrast with synchronous FL, where the server has to wait for all chosen clients to finish their local exercises, in asynchronous FL, the clients can send their updates as soon as they are ready. This can greatly accelerate the training, particularly in heterogeneous environments with straggler clients [Xie et al., 2019]. However, it comes with the issues of staleness of updates and the guarantees of convergence.
- **Federated Dropout and Sub-model Training:** Clients may not train the complete model locally, but a subset of model parameters or sub-model. This results in a lower computation required at the clients and the size (gradient updates) of the communication with the server [see Caldas et al., 2018].

c. Federated TL and RL There are a number of works that train the agent through the FT and the RL.

- **Federated Transfer Learning (FTL):** FTL is a fusion of transfer learning with FL, which transfers knowledge across domains or tasks but does not reveal the data. In transit, this could take the form of pre-training a model on a large public dataset (e.g., general traffic patterns), then fine-tuning with FL using private, transit agency-specific data (e.g., specific bus routes or subway lines) [Liu et al., 2020]. This can be particularly useful in scenarios where some clients have very little labeled data.
- **Federated Reinforcement Learning (FRL):** FRL generalizes the FL framework to the RL scenarios in which, e.g., multiple agents (e.g., autonomous cars, traffic light controllers) aim to collaboratively learn their strategic behavior not jointly sharing their raw environment observations or the reward indications. This is particularly interesting for dynamic systems under transit, like adaptive traffic signal control, autonomous vehicle navigation, or dynamic-pricing ride-sharing, where agents could learn from the system globally, without sacrificing privacy [Nishio & Yonetani, 2019].

d. Privacy-Enhancing Technologies in Transit Contexts

In addition to the natural privacy gain from FL, introducing Extra Privacy- Ensuring Technologies (EPETs) are necessary to enforce strong privacy guarantees, particularly in privacy-sensitive transit scenarios:

- **Homomorphic Encryption (HE):** HE is a way to perform computation on the ciphertext, allowing the centralized server to aggregate encrypted model updates without decrypting them at all. This offers good privacy during aggregation and indeed nothing about individual client's contributions is learned at the server [Hardy et al., 2017]. Computationally expensive as it is, the HE technology development is making it possible for FL under transit.
- **Advanced Differential Privacy Methods:** Using DP in transit settings involves balancing the trade-off between privacy and utility. Recent work addresses adaptive DP mechanisms that only noise should be adjusted according to both data sensitivity and model convergence [Abadi et al., 2016]. Local differential privacy (LDP) can also be implemented at the client side which adds noise before data is transmitted from the end device, providing stronger privacy guarantees at the cost of reduced utility.
- **Tokenized contracts on Trusted Execution Environments (TEEs):** TEEs like Intel SGX or ARM TrustZone, provide an isolated enclave on the processor where code and/or data can be run such that it is guaranteed execution with integrity and confidentiality. In FL, TEEs offer ability to secure the aggregation process on the central server or colocal training selectively on the devices for preserving, not even the central system operator knows about the sensitive model updates or data [Mo et al., 2021].
- **Zero-Knowledge Proofs (ZKPs):** These proofs enable a prover to convince a verifier that some statement holds, without giving them any information other than the fact that the statement is valid. In FL, ZKPs could also be used to verify the correctness of local model updates (e.g., proving that a client has executed the training procedure correctly on its data without disclosing the data or the model) or to comply with certain restrictive usage policies about data usage [Weng et al., 2021].

Together, they are aimed to make FL as a more practical and secure solution for privacy-preserving data analysis in the complex and sensitive setting of urban transit networks.

V. Limitations and Open Challenges

Despite the great achievements and the clear benefits of Federated Learning in preserving privacy for data analysis in transit networks, there still exist several critical limitations and open issues. Solving these issues is essential to achieve the wide acceptance and successful deployment of FL in real-world smart mobility systems.

Although latency and scalability are traditionally treated independently, in practice they are highly coupled in ITS. The problem is how to not only scale federated learning to the order of thousands or even more millions of nodes, but also to make sure such large-scale cooperation also meets time-critical constraints. The lack of existing frameworks that can provide a high degree of scalability with ultra-low latency at the same time, is one of the most urgent gaps for future FL research in transportation.

a. Technical Barriers

- **Latency and real-time constraints:** A large number of transit applications, such as real-time traffic management, dynamic routing and anomaly detection need extremely low latency. The rounds of iterative communication itself in FL, as efficient communication as it may be, may cause delays that are insufficient to meet hard real-time demands. How to perform FL efficiently for ultra-low latency and fast model adaptation in dynamic environments is an open challenge [KonecVný et al., 2016].
- **Scalability in Large Scale Urban Networks:** Urban transport networks consist of a large number of heterogeneous devices and entities, including individual vehicles and smart sensors, as well as multiple transit agencies and other entities. Scaling FL to millions of clients where each may have different computational and network resources is a very hard challenge. Efficient client selection, participation, and aggregation control in these dynamic and large-scale environments while preserving efficiency or

(convergent) is a challenging research field in mathematics [Li et al.,2020].

- **Trust among stakeholders:** Trust among the stakeholders including multiple transit agencies (or their private operators) going into the FL ecosystem is an important while FL offers privacy by design. The for malicious clients injecting poisoned data or model updates to break the integrity of the global model is still an important issue. It is important to build a strong client vetting system, reputation systems, and verifiable computational methods to build trust and reliability in the federated system [Bagdasaryan et al., 2020].
- **Dealing with Extreme Non-IID Data:** While there has been a lot of effort to deal with non-IID data, we still struggle to handle extremely non-IID data, where the data distributions are extremely different across clients, which can lead to significant performance degradation and slow convergence. Designing FL algorithms that are intrinsically robust to highly non-IID data, especially in presence of concept drift or covariate shift, is a key open question [Hsieh et al., 2020].
- **Resource Heterogeneity and Device Constraint:** Clients in a transit FL network can be extremely diverse in terms of their computing resources, storage capacity and power (i.e., edge device mounted on the vehicle vs. powerful server deployed in a data center). This heterogeneity naturally can result in the stragglers where slow clients can slow down the federated learning, which makes it more challenging to design efficient and fair FL algorithms. Balancing FL for resource-tight edge devices without excuse in model performance is a technical challenge [Bonawitz et al., 2019].

b. Governance and Regulation

- **GDPR compliance & data sovereignty:** Operating in a multi-jurisdictional transit landscape under data privacy regulations - such as GDPR (Europe) or CCPA (California) - adds an extra dimension of complexity to FL implementation. Even though FL, as a native consequence of maintaining raw-data local to limit data privacy breaches, provides native mechanisms to satisfy data privacy requirements, which is always a good thing in general no matter whether an absolute compliance can be achieved to complex and evolving data privacy regulations including data minimization, consent management along with retention policy and the right to be forgotten as considerations. Creating FL systems that are native with respect to many regulatory environments is a major open problem [Ziosi et al., 2021].
- **Public Trust and Explainability:** The introduction of AI-driven systems, in particular those processing sensitive information such as mobility patterns, raises public trust concerns. Transparency is crucial for public acceptability: both in how FL models are trained, and in how decisions are made (i.e. explainability). Unexplainability can breed distrust, especially when models are biased or result in unfair decisions. Explainable Federated Learning (XFL) is a critical area of research to mitigate these concerns [Weng et al., 2021].

Cross-Agency Data Collaboration Legal Frameworks: Legal and operational frameworks are necessary for the seamless merging of data across multiple transit agencies and data stakeholders, all with their own data governance policies and competitive interests. Although FL enables collaborations to happen without sharing the raw data, an agreement on who owns the data, who can use the model and who is responsible for the model in a federated setting can be a challenging governance problem.

c. Absence of Benchmark and Datasets in Transport

One of the major obstacles to the development and comparison of FL methods within the domain of transit systems is the lack of publicly available, standardized benchmarks and datasets. The majority of studies using simulation or proprietary databases has restricted for the generalizability and reproducibility of results. Creating more realistic, large scale and diverse datasets that reflect the complexities of urban mobility data (e.g., non-IID nature, real-world noise, multi-modality) is indispensable to ignite novel ideas and to provide a fair ground for various FL algorithms and privacy-enhancing tools. Towards common evaluation metrics and transit use cases definition would also help in advancing research and moving from academic prototypes toward realistic deployments.

Proposed Benchmark Framework for FL in Transport

To address this gap, we propose the creation of standardized benchmarks specifically designed for FL evaluation in transportation contexts. Such benchmarks should encompass the following components:

Table 5. Proposed Benchmark Specifications for FL in Transport

Component	Specification	Example Scenarios
Data Types	Time-series (GPS, sensors), Tabular (tickets, schedules), Images (surveillance), Graph (network topology)	Vehicle trajectories, passenger counts, station cameras, metro line graphs
Data Formats	GTFS (transit feeds), CSV/Parquet (structured), JSON (API responses), HDF5 (large-scale time-series)	Standard GTFS feeds, anonymized ticket logs, IoT sensor streams
Non-IID Scenarios	Label skew, feature skew, quantity skew, temporal distribution shift	Rush hour vs off-peak, urban vs suburban lines, weekday vs weekend patterns
Test Scenarios	Client dropout, adversarial attacks, concept drift, network partitioning	Tunnel connectivity loss, seasonal demand changes, Byzantine clients, multi-agency isolation
Evaluation Metrics	Task accuracy, convergence speed, communication cost, privacy budget, fairness across clients	MAE for demand prediction, F1 for anomaly detection, rounds to convergence, bytes transferred
Scale Levels	Small (10-50 clients), Medium (50-500), Large (500-5000), Massive (5000+)	Single city agencies, metropolitan network, national rail system, global logistics fleet

Such a benchmark framework would enable rigorous comparison of FL algorithms under realistic transport conditions and accelerate the transition from research prototypes to production-ready systems.

VI. Future Directions and Perspectives

The development of Federated Learning in transit networks is on the cusp of substantial growth with continued efforts and growing need for privacy-preserving data analytics. This section discusses important future and emerging directions.

a. Edge Intelligence Integration

The integration of FL with edge computing, which is characterized as Edge AI or Edge Intelligence, will be a key direction in next generation transit systems. By performing computation and model training near the data source, such as in vehicles, smart sensors and roadside units, edge intelligence can greatly reduce communication latency, decrease bandwidth consumption, as well as support real-time decision-making [Zhou et al., 2019]. Future research will focus on:

- **Home FL for Resource-Limited Edge Devices:** Designing lightweight FL algorithms and model architectures which can be efficiently run on computationally, memory, and battery-constrained devices (e.g., transient mobile environments).
- **Dynamic Edge-Cloud Orchestration:** Creating intelligent orchestration mechanisms to dynamically offload FL tasks among edge devices and cloud server according to the network state, data size and computation power, so as to achieve the best performance and resource usage.
- **Secure Edge Aggregation provide security and privacy at edge:** based on Hardware-based TEEs and lightweight cryptographic primitives could be used to protect the local models and updates from tampering or inference attacks at device level.

b. Federated Analytics vs Federated Learning

Though, FL itself is generally used for collaborative training of ML models, a more generic idea, Federated Analytics (FA), has been emerging. FA generalizes the FL approach to support privacy-preserving statistical analysis and aggregation of insights into data directly at participating nodes, without the need for learning a predictive model [Bonawitz et al., 2019]. This difference has implications for transit agencies, which may be more interested in aggregated information (e.g., average travel times, maximum passenger loads, accident rates) for making day-to-day operational planning and policy decisions, than in complex prediction models. Future research will explore:

- **Privacy-Preserving Statistical Aggregation:** To design strong FA mechanisms that can compute arbitrary statistical functions such as sums, averages, counts, histograms across distributed datasets with strong privacy, potentially leveraging stronger differential privacy methods.
- **Federated Query of Analytics:** It leads to transit operators to query the federated network with specific analytical queries and get back aggregated privacy preserving answers which dictates that the data is interpreted in a specific way without revealing the original raw data.
- **Benchmark for FA for Transit:** We aim to define standardized bench- marks and datasets for comparing FA methods in the context of transit, with a focus on estimating the accuracy of the aggregated statistics under different privacy budgets.

c. Cross-Modal FL

Urban transit systems output data of various formats such as videos captured from surveillance cameras, textual reports from online news and social media, numeric reports from sensors and ticketing data, geo-tagged data from GPS transponders. Integrating these heterogeneous sensor data in a privacy-preserving way is essential for comprehensive analysis. Cross-modal FL tries to aggregate information from multi-modal data for a comprehensive and accurate model [Zhang et al., 2021]. Future directions include:

- **Multi-modal Feature Fusion:** Building FL architectures capable of fusing features over distinct data modalities (e.g., visual features from cameras and numerical features from sensors) without spreading the raw data.
- **Federated Multi-task Learning:** Jointly training one FL model for different multiple related tasks performed across different source data types (e.g., predicting traffic congestion and anomaly of sensor and video data, respectively).
- **Privacy-Preserving Cross-Modal Alignment:** Developing methods that can align or correlate data coming from different modalities without disclosing any sensitive information from one modality in particular.

d. Inter-agency FL networks (no exchange of raw data)

The fragmented local governance of the city, and the presence of many independent transit operators serving the same metropolitan area, present an obstacle to the analysis of an extensive data set. Cross-agency FL networks provide an attractive solution that facilitates collaborative intelligence across these agencies while not obligating them to disclose their raw proprietary data. Future research will focus on:

- **Innovations in Collaboration Incentives:** Development of economic and technical incentive structures for a wider range of transit agencies to establish a more open participation structure that eases concerns about gaining a competitive advantage or sharing resources.
- **Standardized FL Protocols for Inter-Agency Use:** Establishing standard protocols, exchange data format and common API for inter-agency use so that FL system can be easily integrated, between departments and their disparate IT infrastructure.
- **Legal and Policy Frameworks:** Work with legal professionals and policy makers to develop explicit legal and policy frameworks that will promote inter-agency FL networks and resolve issues of data ownership, liability, and regulatory adherence in a multi-jurisdictional environment.

These future directions suggest that FL research in transit is a dynamic and emergent field. Focusing on these directions, FL can really reinforce the position of a founding technology for the development of intelligent, privacy-preserving, and resilient urban mobility systems.

From the limitations, scalability with real time constraints is one of the most serious issues for future ITS applications. In this thesis we will specifically address this gap by presenting scalable federated learning approaches that can support the operation of large-scale urban mobility systems satisfying tight latency constraints.

VII. Conclusion

This paper has presented an in-depth survey of advancements in the development of federated learning approaches towards privacy-preserving data analytics in transit networks. We started by emphasizing the importance of data analytics in the context of modern smart transit systems and articulated the privacy threats that arise from centralized-level data processing in these systems. Federated learning, a novel paradigm, can tackle the above issues by allowing collaborative model training while avoiding the sharing of raw sensitive data.

We discussed the fundamental concepts of FL such as federated averaging, secure aggregation, differential privacy, and homomorphic encryption, which in combination enhance the privacy and security guarantees of the framework. In addition, the applications of FL in different transit scenarios were discussed including real-time passenger analytics, fare evasion detection, predictive maintenance, traffic flow management, and anomaly detection, which highlighted its flexibility and the potential use of the practical impact. Of course, we have to critically inspect the emerging challenges of deployment including non-IID data, communication overhead, real-time constraints, system heterogeneity, data quality and adversarial attack issues.

The most recent developments of FL research have been addressing such limitations through the personalization of FL, communication-efficient algorithms, federated transfer and reinforcement learning and the integration of advanced privacy-ensuring technologies such as Trusted Execution Environments and Zero-Knowledge Proof. In the future of FL in transit, we envisage the promising opportunities including the integration of edge intelligence, federated analytics, cross-modal FL, as well as the construction of resilient inter-agency FL networks. Such directions are expected to improve the efficiency, scalability, and privacy guarantees of FL solutions.

In summary, federated learning is a highly disruptive technology for urban mobility, as it provides the perspective for the understanding of mobility data in an effective manner, without compromising the privacy of the individuals and organizations. Realizing this potential will require further cross-disciplinary work among investigators, transit providers, policy makers, and industry. By tackling together, the remaining technical, governance, and ethical challenges we can open up a new era of intelligent, secure, and privacy- preserving smart cities, where data-driven insights will seamlessly improve urban mobility for all citizens.

Despite the great efforts so far dedicated to the application of Federated Learning (FL) to the intelligent transportation systems, current solutions are still not able to ensure at the same time the scalability on millions of diverse devices, which the automotive scenario envisages, and the ultra-low latency, which is necessary for RT applications, such as traffic management, anomaly detection, or predictive maintenance. The existing works either trade off privacy with efficiency, or optimize convergence but are effective only on extremely non-IID data and resource heterogeneity scenarios. Moreover, due to the absence of benchmark datasets for transportation scenarios, it is difficult for people to compare achievements and promote real-world deployments. This open research problem of being scalable yet having low latency is a major gap in the literature. It is this gap that the thesis seeks to address, by designing and evaluating privacy-preserving, robust and scalable FL frameworks which are envisioned to serve large-scale, real-time transit network

References

[1] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Federated learning of deep networks using model averaging. 20th International Conference on Artificial Intelligence and Statistics (AISTATS).

[2] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patrick, J., Talwar, K., Trevithick, L., & Vassilvitskii, S. (2017). Securing the aggregation in federated learning. In: CCS 17: 2017 ACM SIGSAC Conference on Computer and Communications Security.

[3] Dwork, C., Roth, A., et al. (2014). Theoretical foundations for differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.

[4] Gentry, C. (2009). Fully Homomorphic Encryption from Ideal Lattices. In Proceedings of the 41st Annual ACM Symposium on Theory of Computing.

[5] Zhang, S., Li, J., Shi, L., Ding, M., Nguyen, D. C., & Dobre, O. A. (2023). Federated learning in ITS: Recent applications and open issues. IEEE Transactions on Intelligent Transportation Systems

[6] Chong, Y. W., Yau, K. L. A., Ibrahim, N. F., & Han, J. (2024. Federated Learning for Intelligent Transport Systems: Use Cases, Open Challenges, and Opportunities. IEEE Transactions on Intelligent Transportation Systems.

[7] Wang, S., Hu, T., & Min, G. (n.d.). Federated learning in intelligent transportation systems: A survey. IEEE Transactions on ITS.

[8] Gupta, A., et al. (2020). Blockchain and AI in fare collection in public transport. Journal of Urban Mobility.

[9] Chen, Y., & Zhang, L. (2021). Federated learning approach for predictive maintenance in smart transportation systems. Journal of Intelligent Transportation Systems.

[10 Lu, Y., et al. (2021). Traffic flow prediction in smart cities based on federated learning. IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS.

[11] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challanges, methods and future directions

[12] Konečný, J. McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016).
Federated learning: Towards better communication efficiency. arXiv preprint arXiv:1610.05492.

[13] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, F., Ingerman, A., Jackson, N., … & Van Der Maaten, L.ecru (2019). Towards vigorous federated learning at scale: System design. In Proceedings of Machine Learning and Systems, 1, 371–382.

[14] Bagdasaryan, E., Shen, A., Sharma, A., Fong, M., Veeramachaneni, V., Sinha, N., & Shmatikov, V. (2020). How to backdoor federated learning. International Conference on Artificial Intelligence and Statistics.

[15] Fallah, A., Mokhtari, A., & Ozdaglar, A. (2020). Towards practical federated learning: fairly trading on clients' data. In Advances in Neural Information Processing Systems 33.

[16] Acar, D. A. E., Malek, A., & Kairouz, P. (2021). Federated Learning with Heterogenous Clients: A Unified View. arXiv preprint arXiv:2102.08092.

[17] Cho, Y. J., & Kim, J. (2020). Client sampling for federated learning with non-IID data in mobile edge networks. IEEE Access, 8, 195759–195769.

[18] Wang, S., Wang, T., & Liu, J. (2020). Adaptive federated learning via dot product transfer. In AAAI conference on artificial intelligence (Vol. 34, No. 04, pp. 6296-6303).

[19] Alistarh, D., Grubic, D., Li, J., Tomioka, R., \& Vojnovic, M. (2017). Qsparse-local-SGD: Decentralized SGD with Quantization and Sparsification and Local Computations. In Neural Information Processing Systems, NIPS, 30.

[20] Xie, C., Koyejo, S., and Gupta, I. (2019). Asynchronous federated optimization. arXiv preprint arXiv:1903.03934.

[21] Caldas, A., Duddu, S. M. R., Li, P., Konečný, J., McMahan, H. B., & Smith, V. (2018). Extending the scale of federated learning via reducing client computational burden. arXiv preprint arXiv:1812.07279.

[22] Liu, Y., Kang, J., & Niyato, D. (2020). Vehicular networks federated transfer learning. IEEE Wireless Communications Letters, 9(10), 1730–1734.

[23] Nishio, T., & Yonetani, R. (2019).(There's one another method by which ANNs can form specific expectations (e.g. by training, He et al., [2015a]).

[24] Hardy, S., Henecka, M., Pfitzmann, A., & Waidner, M. (2017). Private federated learning on the plaintext. arXiv preprint arXiv:1711.08060.

[25] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G., Davis, A., Dean, J., Devin, M., & Ghemawat, S. (2016) TensorFlow: Large-scale machine learning on heterogeneous systems. Deep learning with differential privacy. In ACM Computer Communication Security (CCS).

[26] Mo, F., Yu, Y., Li, X., & Wu, X. (2021) Private federated learning under a malicious but honest but curious model with a trust execution environment. Proceedings of the IEEE International Conference on Computer Design.

[27] Weng, J., Zhang, S., & Li, J. (2021). Towards explainable federated learning. arXiv preprint arXiv:2103.01894.

[28] Ziosi, M., et al. (2021). Bias and privacy in Ai-driven public transport surveillance: A framework for responsible innovation. Ethics and Information Technology.

[29] Zhou, Z., Chen, X., Li, J. (2019). Edge intelligence: The next wave of AI. Proceedings of the IEEE, 107(8), 1738–1762.

[30] Zhang, S., Li, J., and Shi, L. Federated multi-view learning for intelligent transportation systems. IEEE Transactions on ITS.

[31] Hsieh, K., &Yang, J. (2020). The challenges of non-IID data in federated learning: A survey. arXiv preprint arXiv:2007.01439.