

Federated Learning for Privacy-Preserving Vehicle Trajectory Prediction in Intelligent Transportation Networks

Kirsty Hyde

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.

hydekirsty@oregonstate.edu

Oran Skehan

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.

oranmail@unr.edu

Abstract

Vehicle trajectory prediction has become a foundational capability for intelligent transportation networks, supporting collision avoidance, cooperative routing, signal optimization, and mixed-autonomy traffic management. Conventional centralized prediction models require large-scale movement data to be aggregated in cloud repositories, creating substantial privacy, cybersecurity, and governance risks because mobility traces can reveal home locations, work patterns, health visits, social ties, and other sensitive inferences. Federated learning offers a distributed alternative in which vehicles, roadside units, and edge servers collaboratively train prediction models without transmitting raw trajectories to a central authority. This paper examines federated learning for privacy-preserving vehicle trajectory prediction from a system-level perspective. It argues that privacy protection cannot be reduced to a single technical mechanism, but must be understood as an architectural and institutional property emerging from the interaction of secure aggregation, differential privacy, client sampling, communication scheduling, edge computing, and legal accountability. The paper analyzes structural trade-offs involving prediction accuracy, personalization, communication cost, latency, energy use, robustness, fairness, and sustainability. It further considers governance challenges in intelligent transportation ecosystems, including consent, purpose limitation, auditability, liability, and cross-jurisdictional data policy. Through conceptual analysis, cross-domain comparisons, and forward-looking deployment scenarios, the paper shows that federated trajectory prediction is not merely a privacy-enhancing technique but a socio-technical infrastructure requiring coordinated design across models, networks, institutions, and public policy. The conclusion outlines research and policy priorities for deploying federated prediction systems that are accurate, equitable, resilient, and democratically accountable.

Keywords

federated learning; vehicle trajectory prediction; intelligent transportation systems; privacy; edge computing; governance; robustness; fairness; sustainability; data policy.

1. Introduction

Intelligent transportation networks increasingly depend on predictive intelligence rather than reactive control. Vehicles, roadside sensing units, mobile devices, and urban management platforms generate continuous streams of positional, kinematic, and contextual data. When these streams are processed through machine learning models, they can support short-term trajectory forecasting, intersection safety, lane-change anticipation, pedestrian interaction modeling, and cooperative maneuver planning. The academic foundations of trajectory prediction have evolved through social pooling, sequence modeling, generative adversarial modeling, and increasingly complex interaction-aware representations that seek to capture the mutual influence among vehicles, cyclists, pedestrians, and infrastructure [13] [14]. In parallel, urban computing has demonstrated that mobility data can improve transportation efficiency, environmental management, and public service planning, but also that such data carry profound privacy and equity implications [20]. The central tension is therefore not whether prediction is useful, but how predictive capability can be achieved without reproducing the surveillance-intensive architectures that centralized data aggregation tends to encourage.

Federated learning has emerged as a leading response to this tension. Instead of moving raw data to a model, federated learning moves model updates to a coordinating server or uses decentralized aggregation protocols so that training occurs locally on vehicles, roadside units, or edge nodes [1] [2]. The promise is considerable: sensitive trajectories remain closer to their point of origin, communication can be reduced through local computation, and model improvement can continue across fleets without direct data pooling. However, the system-level reality is more complicated. Vehicle trajectory prediction is not a generic classification task. It involves sequential dependencies, multi-agent interactions, heterogeneous sensing quality, rapidly changing environments, and safety-critical latency constraints. A federated architecture that performs well in static, well-connected, and homogeneous environments may fail in sparse rural corridors, dense urban intersections, or mixed fleets with uneven computational resources. The design problem is therefore architectural, not merely algorithmic.

This paper takes a broad socio-technical view of federated learning for privacy-preserving vehicle trajectory prediction in intelligent transportation networks. It emphasizes structural trade-offs, infrastructure dependencies, governance arrangements, deployment pathways, sustainability, robustness, and fairness. The argument developed here is that federated learning should not be treated as a privacy switch that can be attached to an existing prediction pipeline. Rather, it should be understood as a distributed system whose privacy guarantees, predictive performance, and social legitimacy depend on choices about aggregation, client participation, network topology, edge capacity, regulatory oversight, and public accountability. These choices determine whether federated trajectory prediction becomes a trustworthy transportation infrastructure or a fragmented collection of technically sophisticated but institutionally fragile pilots.

2. Privacy and Data Governance in Intelligent Transportation Networks

Privacy in vehicle trajectory prediction is not limited to hiding vehicle identifiers. Even pseudonymized mobility traces can be re-identified through spatiotemporal uniqueness, route regularity, and linkage with external datasets. A trajectory that begins near a residential neighborhood, passes through a medical district, and ends near a workplace can disclose sensitive information about health, employment, religion, or social relationships. Machine learning models trained on such data can also leak information through membership inference, model inversion, and reconstruction attacks [7] [8]. These risks are especially acute in

transportation because mobility is continuous, relational, and difficult to anonymize without destroying much of its predictive value. Federated learning reduces the exposure of raw data, but it does not by itself eliminate inference risk, because model updates and gradients can still encode sensitive information.

Secure aggregation provides an important architectural layer by ensuring that a central server cannot inspect individual client updates during training [4]. Differential privacy adds a formal statistical guarantee by limiting the influence of any single record or client on the learned model [5] [6]. When combined, these mechanisms can provide stronger privacy than either approach alone. Yet the combination also introduces difficult system trade-offs. Strong differential privacy noise can degrade prediction accuracy, especially for rare but safety-critical events such as unusual pedestrian movements or atypical vehicle interactions. Secure aggregation can increase communication overhead, complicate client dropout handling, and reduce the ability to diagnose malicious or corrupted updates. Privacy is therefore not a free property; it is a budget that must be allocated across training rounds, model layers, client groups, and operational contexts.

Governance is equally important. Privacy-preserving federated learning requires clear rules about what data may be used, for what purposes, for how long, and under whose authority. In transportation networks, stakeholders include vehicle manufacturers, fleet operators, public agencies, infrastructure owners, software vendors, insurers, and travelers themselves. Each has different incentives and legal obligations. A federated system that complies with one jurisdiction may violate another, particularly when model updates cross national borders or when public agencies seek access for enforcement or planning. Purpose limitation, data minimization, retention limits, auditability, and independent oversight are therefore not bureaucratic additions to federated learning; they are constitutive elements of trustworthy deployment. Recent work on high-density traffic and diverse agent-interaction scenarios [15] highlights that prediction models must handle complex social and environmental contexts, which in turn increases the need for careful governance because richer interaction data can reveal more about individual behavior.

A further governance challenge concerns consent and collective harm. Individual drivers may consent to participate in a federated training program, but their data may also contribute to models that affect pedestrians, cyclists, and communities that did not consent. Trajectory prediction can shape policing, insurance pricing, infrastructure investment, and access to public space. Privacy must therefore be understood as a collective and distributive concern, not only an individual choice. Federated learning can support more localized control, but it can also fragment accountability if no single actor is responsible for the overall system. Effective governance requires mechanisms for transparency, redress, impact assessment, and participation by affected communities, especially where transportation systems have historically produced unequal outcomes.

3. Federated Learning as a System Architecture for Trajectory Prediction

Federated learning architectures for vehicle trajectory prediction can be organized around several structural decisions. The first concerns the unit of federation. Clients may be individual vehicles, groups of vehicles within a geographic zone, roadside units, edge data centers, or municipal transportation authorities. Vehicle-level federation maximizes data locality but creates severe challenges related to churn, connectivity, and heterogeneous computation. Roadside-unit federation can aggregate local observations and reduce uplink traffic, but it may concentrate data in infrastructure that is itself vulnerable to physical and

cyber threats. Municipal federation can align with public accountability but may exclude private fleets and consumer vehicles. Each choice redistributes privacy risk, operational control, and system complexity.

The second decision concerns aggregation. Centralized federated averaging remains the most widely studied approach, but it depends on a coordinating server and may struggle with non-identically distributed data [2] [9]. Vehicles in different regions, road types, and traffic cultures generate different trajectory distributions. A model trained on dense urban traffic may perform poorly in rural highways, while a model trained on passenger vehicles may not generalize to trucks, buses, motorcycles, or emergency vehicles. Techniques such as matched averaging and stochastic controlled averaging attempt to address client drift and improve convergence under heterogeneity [11] [12]. Personalized federated learning goes further by allowing local adaptation, but personalization can weaken privacy if local models become too closely fitted to individual trajectories. The design challenge is to balance global generalization with local relevance without turning the federated model into a de facto repository of individualized behavior.

The third decision concerns communication and synchronization. Trajectory prediction models often use recurrent, convolutional, graph-based, or transformer-like sequence representations that can be large and computationally demanding. Frequent synchronization may be impractical in vehicular networks where connections are intermittent and bandwidth is shared with safety-critical messages. Asynchronous aggregation, periodic updates, gradient compression, quantization, and event-triggered participation can reduce communication load, but they can also introduce staleness and bias [3]. In safety-critical settings, stale models may produce unreliable predictions precisely when traffic conditions change rapidly. Communication efficiency must therefore be evaluated not only in terms of bytes saved but also in terms of prediction reliability under delay, packet loss, and partial participation.

The fourth decision concerns the relationship between prediction and privacy. Federated learning is often combined with differential privacy, secure aggregation, or both. However, the privacy guarantee for trajectory prediction depends on the sensitivity of the update, the number of participating clients, the sampling strategy, and the temporal correlation of mobility data. A single vehicle's updates over many rounds can reveal more than a single snapshot, and repeated participation can amplify privacy loss. Privacy accounting must therefore be longitudinal rather than per-round. At the same time, the model must remain accurate enough for downstream safety functions. This tension is particularly difficult because rare events, such as near-collisions or unusual pedestrian interactions, are often the most valuable for training and the most sensitive for privacy. The system must decide how to represent rare events without singling out the individuals or vehicles involved.

Federated trajectory prediction also requires careful treatment of evaluation. Centralized accuracy metrics may hide disparities across regions, vehicle types, and vulnerable road users. A model that achieves high average displacement error may still fail at night, in rain, or in low-income neighborhoods with poor infrastructure. Robust evaluation must include subgroup performance, calibration, failure cases, and adversarial conditions. Recent advances in generating high-density and diverse interaction scenarios [15] can support more rigorous stress testing, but synthetic scenarios cannot replace real-world validation across social and geographic contexts. Evaluation should therefore be treated as a governance instrument, not only an engineering step, because it determines which failures become visible and which remain hidden.

4. Infrastructure, Communication, and Deployment Trade-offs

The deployment of federated trajectory prediction depends on a layered infrastructure of onboard computers, roadside units, edge servers, cellular and dedicated short-range communication links, and cloud-based orchestration. Each layer has different latency, reliability, energy, and security characteristics. Onboard computation preserves privacy and reduces communication but is constrained by vehicle power, thermal limits, and hardware heterogeneity. Edge servers can aggregate local updates, provide lower latency than distant clouds, and support regional personalization, but they require trust, physical security, and operational maintenance. Cloud orchestration can coordinate large fleets and manage global model versions, but it reintroduces central points of control and potential surveillance. A resilient architecture is likely hybrid, with latency-sensitive functions near the vehicle and slower model improvement functions distributed across regional and central tiers.

Vehicular networks introduce distinctive communication constraints. Vehicle-to-everything communication must support safety applications with very low latency and high reliability, while federated learning traffic competes for the same spectrum and backhaul capacity [17]. If model updates are large or frequent, they may degrade safety messaging or increase cost for vehicle owners and municipalities. Edge caching, model compression, and participatory scheduling can mitigate these effects, but they also create new fairness questions. Vehicles with better connectivity and newer hardware may contribute more frequently, while older vehicles, rural users, and low-income communities may be underrepresented. That underrepresentation can bias the global model and reproduce existing transportation inequalities. Federated learning for vehicular Internet of Things therefore requires network-aware client selection and resource allocation that explicitly consider equity, not only throughput [16].

Security is another deployment concern. Federated systems are vulnerable to poisoning, Sybil attacks, backdoors, and free-riding. An adversary could inject corrupted updates that degrade prediction in specific locations or trigger unsafe behavior under particular conditions. Blockchain-based coordination and distributed ledgers have been proposed to improve transparency and trust in federated data sharing, including industrial Internet of Things settings [18]. However, blockchain introduces its own scalability, energy, and governance challenges. Similarly, federated anomaly detection systems in IoT demonstrate the value of decentralized learning for identifying unusual behavior without centralizing raw data [19], but they also show that detection quality depends on diverse and representative participation. In transportation, a poisoned model could affect not only privacy but physical safety, making robust aggregation, update validation, and human oversight essential.

Sustainability must be considered alongside performance. Federated learning can reduce data transmission and centralized computation, but it shifts energy consumption to vehicles and edge devices. Frequent training on onboard hardware may increase battery drain in electric vehicles, reduce vehicle range, and shorten device lifespan. Edge servers require continuous power and cooling. Communication networks consume energy for transmission and routing. The net environmental effect depends on model size, update frequency, network efficiency, and the carbon intensity of the electricity grid. A sustainable design would use adaptive participation, event-triggered training, model reuse, and regional scheduling that aligns computation with renewable energy availability. Sustainability is therefore not an afterthought but a design constraint that interacts with privacy, latency, and equity.

5. Robustness, Fairness, and Sustainability

Robustness in federated trajectory prediction has multiple dimensions. Statistical robustness concerns the ability to learn from non-identical and non-stationary data [9] [10]. Operational robustness concerns resilience to communication failures, sensor degradation, cyberattacks, and model drift. Safety robustness concerns the consequences of prediction errors in physical environments. A model that is robust on average may still fail catastrophically in rare but high-stakes situations. Federated learning can improve robustness by pooling knowledge across diverse environments, but it can also amplify vulnerabilities if malicious or biased clients dominate aggregation. Robust aggregation rules, anomaly detection, and trust scoring can help, but they must be designed to avoid excluding legitimate minority populations. The challenge is to distinguish between harmful outliers and valuable rare cases, especially when rare cases involve vulnerable road users or uncommon traffic conditions.

Fairness is closely related. Federated systems can either reduce or intensify disparities depending on how clients are selected, how models are aggregated, and how performance is evaluated. If training prioritizes dense urban areas with abundant data and reliable connectivity, rural and underserved regions may receive less accurate predictions. If personalization is offered only to premium vehicles or fleet customers, safety benefits may become unequally distributed. If privacy budgets are allocated uniformly, small groups may receive weaker protection because their updates are more identifiable. Fair federated learning therefore requires explicit attention to representation, allocation, and harm. It also requires participatory governance, because technical definitions of fairness may not align with community priorities. In transportation, fairness intersects with accessibility, environmental justice, and the distribution of public space.

Sustainability and fairness are linked through infrastructure. Regions with less reliable power, fewer edge servers, and older vehicles may struggle to participate in federated training, yet they may need accurate prediction most. A federated architecture that assumes continuous connectivity and modern onboard computation will systematically exclude those regions. Lightweight models, intermittent participation, and cross-silo aggregation through public agencies or community organizations can broaden access. However, these solutions require investment and coordination. Without deliberate policy support, federated learning may become a premium capability for well-resourced mobility ecosystems while under-resourced areas remain dependent on less accurate or less private centralized systems. The socio-technical design must therefore treat infrastructure equity as a first-class objective.

Privacy itself has a fairness dimension. Differential privacy noise can reduce accuracy for small subgroups more than for large ones. Secure aggregation can make it harder to audit whether a model is biased against a particular group. The privacy-utility trade-off is not experienced uniformly. A model that is sufficiently private for a large fleet may still reveal information about a single vehicle in a sparsely populated region. Conversely, strong privacy protections may reduce the model's ability to detect rare but dangerous events affecting pedestrians or cyclists. These trade-offs cannot be resolved by technical optimization alone. They require public deliberation about acceptable risk, accountability, and the distribution of benefits and burdens. Recent work on high-density and diverse interaction scenarios [15] can help evaluate model behavior under complex conditions, but it also underscores that prediction quality is inseparable from the social contexts in which vehicles operate.

6. Policy, Governance, and Cross-Domain Comparisons

Policy frameworks for federated trajectory prediction must address data protection, cybersecurity, transportation safety, competition, and public procurement. Data protection

laws often emphasize lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity, and accountability. Federated learning can support data minimization by keeping raw trajectories local, but it does not automatically satisfy transparency or accountability. Regulators may need access to model documentation, privacy audits, impact assessments, and performance reports. Secure aggregation and differential privacy can provide technical evidence, but they must be translated into institutional assurances that non-experts can evaluate [4] [6]. Without such translation, federated systems may suffer from a legitimacy deficit even when their technical privacy properties are strong.

Cross-domain comparisons offer useful lessons. In healthcare, federated learning is used to train models across hospitals without sharing patient records, but it faces challenges of institutional heterogeneity, regulatory fragmentation, and unequal data quality. In finance, federated learning can support fraud detection across institutions while respecting secrecy, but it raises concerns about collusion, market power, and consumer protection. In industrial IoT, federated learning and blockchain have been combined to support privacy-preserving data sharing and anomaly detection [18] [19]. These domains show that federated systems succeed when trust is distributed, incentives are aligned, and governance is embedded in the architecture. They also show that technical privacy alone is insufficient when power asymmetries allow a dominant actor to infer sensitive information from model updates or metadata.

Transportation adds a public infrastructure dimension. Roads, signals, and mobility data are often governed by public agencies, yet much of the technology is developed by private firms. Federated learning can distribute model training, but it does not resolve questions about who owns the global model, who can update it, who can audit it, and who is liable when it fails. A public agency may require transparency and safety validation, while a manufacturer may claim proprietary interests. A fleet operator may want personalized routing, while a city may prioritize congestion reduction and equity. Governance must therefore include contracts, standards, procurement rules, and oversight bodies. It must also address cross-jurisdictional operation, because vehicles routinely cross municipal, state, and national borders. Model updates that are lawful in one region may be restricted in another, creating compliance complexity and potential fragmentation.

Policy should also anticipate dual-use risks. Trajectory prediction can improve safety, but it can also support surveillance, predictive policing, targeted advertising, and discriminatory pricing. Federated learning may reduce raw data centralization, but the resulting models can still be used for purposes beyond the original consent. Purpose limitation and use restrictions are therefore essential. Technical mechanisms such as secure enclaves, access control, and model watermarking can help, but they cannot replace legal and institutional constraints. Public participation, independent review, and whistleblower protection are important for detecting mission creep. In this sense, federated learning is not inherently democratizing; its social effects depend on the governance regime in which it is embedded.

7. Case Illustrations and Forward-Looking Perspectives

Consider a smart corridor with mixed autonomy, pedestrians, cyclists, buses, and delivery vehicles. A federated prediction system could train on vehicles and roadside units within the corridor, using secure aggregation to protect individual traces and differential privacy to limit inference. Edge servers at intersections could personalize models for local geometry and traffic patterns, while a regional coordinator maintains a global model for rare events and

cross-corridor consistency. During peak hours, communication may be constrained, so only a subset of vehicles participate. During off-peak hours, larger updates can be transmitted. The system could use synthetic high-density scenarios to stress-test rare interactions and then validate them with real-world observations [15]. However, the same system could be misused to track specific vehicles if update metadata, timing, and location are not protected. The case illustrates that privacy depends on the entire data flow, not only the training algorithm.

A second case involves rural highways and long-haul freight. Here, connectivity may be intermittent, vehicles may be large and computationally capable, and traffic density is low. Federated learning could allow trucks to share knowledge about weather, road conditions, and fatigue-related driving patterns without exposing proprietary routes or driver behavior. But low-density participation increases privacy risk because each client's update is more identifiable. It also increases model bias because the data distribution is narrow. A hybrid approach might combine federated learning with local personalization, public weather data, and infrastructure sensors. Governance would need to address the interests of drivers, carriers, insurers, and regulators. The case shows that federated learning must adapt to geographic and economic context rather than assuming a uniform urban model.

A third case concerns emergency response. Ambulances, fire trucks, and police vehicles need accurate trajectory prediction under unusual conditions, including temporary road closures, crowd movement, and rapid route changes. A federated model could learn from emergency vehicle trajectories without centralizing sensitive operational data. However, adversarial actors might try to infer emergency patterns or disrupt prediction during critical events. Robustness, access control, and fail-safe fallback modes become paramount. The system must be able to operate safely when the federated network is degraded or under attack. This requires redundancy, human oversight, and clear liability rules. It also requires that privacy protections not prevent legitimate after-action review and public accountability.

Forward-looking perspectives point toward several research and policy priorities. First, federated trajectory prediction needs longitudinal privacy accounting that captures repeated participation and temporal correlation. Second, it needs robust and fair aggregation methods that protect minority groups and rare events. Third, it needs infrastructure models that account for energy, bandwidth, maintenance, and geographic equity. Fourth, it needs governance frameworks that combine technical audits with public participation and legal accountability. Fifth, it needs standards for interoperability so that models and privacy mechanisms can work across manufacturers, municipalities, and jurisdictions. These priorities are interconnected. A technical advance in privacy may shift the governance burden; a policy change may alter the feasible architecture. The field should therefore be pursued as a transdisciplinary program rather than a narrow machine learning problem.

8. Conclusion

Federated learning offers a credible architectural path toward privacy-preserving vehicle trajectory prediction in intelligent transportation networks. By keeping raw mobility data closer to its source and sharing model updates rather than trajectories, federated systems can reduce central surveillance risks and support collaborative learning across fleets, infrastructure, and regions. However, privacy is not achieved by federation alone. It requires secure aggregation, differential privacy, robust client selection, longitudinal privacy accounting, and governance mechanisms that constrain purpose, access, and reuse. It also requires attention to accuracy, personalization, communication cost, latency, energy,

robustness, fairness, and sustainability. These dimensions interact, and optimizing one can undermine another.

The system-level argument of this paper is that federated trajectory prediction should be designed as a socio-technical infrastructure. Technical choices about aggregation, edge placement, and model compression shape who can participate and who benefits. Governance choices about consent, auditability, liability, and public oversight shape whether the system is legitimate and accountable. Infrastructure choices about connectivity, power, and maintenance shape whether the system is equitable and sustainable. If these dimensions are treated separately, federated learning may become a collection of impressive pilots that fail to scale or earn public trust. If they are treated together, federated learning can support transportation systems that are safer, more efficient, and more respectful of civil liberties. The challenge is not simply to train better models without centralizing data. It is to build a distributed, governable, and equitable prediction infrastructure for the mobility systems of the future.

References

1. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210. <https://doi.org/10.1561/22000000083>
2. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282). PMLR.
3. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
4. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175–1191). ACM. <https://doi.org/10.1145/3133956.3133982>
5. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318). ACM. <https://doi.org/10.1145/2976749.2978318>
6. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
7. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy* (pp. 3–18). IEEE. <https://doi.org/10.1109/SP.2017.41>

8. Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1322–1333). ACM. <https://doi.org/10.1145/2810103.2813677>
9. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. *arXiv preprint arXiv:1806.00582*.
10. Hsu, T.-M. H., Qi, H., & Brown, M. (2019). Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*.
11. Wang, H., Yurochkin, M., Sun, Y., Papailiopoulos, D., & Khazaeni, Y. (2020). Federated learning with matched averaging. In *International Conference on Learning Representations*. ICLR.
12. Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S. J., Stich, S. U., & Suresh, A. T. (2020). SCAFFOLD: Stochastic controlled averaging for federated learning. In *Proceedings of the 37th International Conference on Machine Learning* (pp. 5132–5143). PMLR.
13. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., & Savarese, S. (2016). Social LSTM: Human trajectory prediction in crowded spaces. In *2016 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 961–971). IEEE. <https://doi.org/10.1109/CVPR.2016.110>
14. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., & Alahi, A. (2018). Social GAN: Socially acceptable trajectories with generative adversarial networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2255–2264). IEEE. <https://doi.org/10.1109/CVPR.2018.00240>
15. Yang, R., Xu, Y., Chen, Y., Fu, Y., & Su, L. (2025). A Trajectory Generator for High-Density Traffic and Diverse Agent-Interaction Scenarios. *arXiv preprint arXiv:2510.02627*.
16. Du, Z., Wu, C., Yoshinaga, T., Ji, Y., & Zhang, Y. (2020). Federated learning for vehicular internet of things: Recent advances and open challenges. *IEEE Open Journal of the Computer Society*, 1, 45–61. <https://doi.org/10.1109/OJCS.2020.3003366>
17. Ye, H., Liang, L., Li, G. Y., Kim, J., & Poor, H. V. (2020). Federated learning in vehicular edge computing: Opportunities and challenges. *IEEE Network*, 34(6), 38–45. <https://doi.org/10.1109/MNET.011.2000086>
18. Lu, Y., Huang, X., Zhang, K., Maharjan, S., & Zhang, Y. (2020). Blockchain and federated learning for privacy-preserved data sharing in industrial IoT. *IEEE Transactions on Industrial Informatics*, 16(6), 4177–4186. <https://doi.org/10.1109/TII.2019.2942823>
19. Nguyen, T. D., Marchal, S., Miettinen, M., Fereidooni, H., Asokan, N., & Sadeghi, A.-R. (2019). D²IoT: A federated self-learning anomaly detection system for IoT. In *2019 IEEE 39th International Conference on Distributed Computing Systems* (pp. 756–767). IEEE. <https://doi.org/10.1109/ICDCS.2019.00080>
20. Zheng, Y., Capra, L., Wolfson, O., & Yang, H. (2014). Urban computing: Concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology*, 5(3), 1–55. <https://doi.org/10.1145/2629592>