Federated Multi-Generator Spatiotemporal Modeling for Privacy-Preserving Autonomous Navigation Prediction

Sagar Arora
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
sagar.arora616@uab.edu

Felix Day
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
day1999@ucf.edu

Abstract

The rapid deployment of autonomous vehicles and intelligent navigation systems hinges on the ability to predict future trajectories of dynamic agents with high accuracy while preserving the privacy of individual data contributors. Conventional centralized spatiotemporal models require the aggregation of sensitive localization and behavioral data, raising critical concerns regarding data sovereignty, regulatory compliance, and adversarial vulnerability. This paper proposes a federated multi-generator spatiotemporal modeling paradigm that integrates multiple generative trajectory predictors under a federated learning framework to enable privacy-preserving autonomous navigation prediction. The architecture distributes model training across local nodes, each operating on private data, and aggregates only model parameters or synthetic representations rather than raw trajectories. We examine structural trade-offs between prediction fidelity, communication efficiency, and privacy guarantees across various federation topologies. The multi-generator design leverages ensemble diversity to capture complex spatiotemporal dependencies while also providing robustness against distributional shifts and adversarial attacks. Governance mechanisms such as differential privacy budgets, secure aggregation protocols, and decentralized audit trails are analyzed in the context of real-world deployment constraints. Infrastructure considerations including computational heterogeneity, bandwidth limitations, and latency requirements are discussed alongside sustainability and fairness implications. Through cross-domain comparisons with centralized, single-generator, and non-federated approaches, we highlight the advantages and limitations of the proposed framework. The paper concludes with a forward-looking perspective on policy frameworks, standardization efforts, and ethical guidelines necessary for the responsible adoption of federated multi-generator systems in autonomous navigation.

Keywords

Federated learning, spatiotemporal modeling, multi-generator prediction, autonomous navigation, privacy preservation, trajectory prediction, decentralized infrastructure, governance.

1. Introduction

Autonomous navigation systems rely on accurate and timely prediction of the future positions of surrounding vehicles, pedestrians, cyclists, and other agents to plan safe and efficient trajectories. Spatiotemporal models, particularly those based on graph neural networks and recurrent architectures, have demonstrated state-of-the-art performance in capturing the complex interactions and motion patterns inherent in traffic scenes [2, 6]. However, these models are typically trained in a centralized manner, requiring the aggregation of vast amounts of trajectory data collected from numerous vehicles and infrastructure sensors. This centralization exposes sensitive information, such as location histories, driving habits, and behavioral patterns, to potential breaches, misuse, or unauthorized surveillance. Regulatory frameworks such as the General Data Protection Regulation and the California Consumer Privacy Act impose strict requirements on the collection and processing of personal data, thereby motivating the development of privacy-preserving alternatives.

Federated learning has emerged as a promising paradigm for collaborative model training without centralizing raw data [1, 4]. In federated settings, each client device trains a local model on its own data and shares only model updates, such as gradients or weight parameters, with a central server. While this reduces the exposure of raw data, it does not fully eliminate privacy risks, as model updates can still leak information about the training data through gradient inversion or membership inference attacks. To address these vulnerabilities, recent works have integrated differential privacy mechanisms into federated learning pipelines, adding calibrated noise to model updates to bound the information leakage [5, 8]. Simultaneously, spatiotemporal trajectory prediction has benefited from multi-generator architectures that combine multiple predictive models, each capturing different aspects of motion dynamics, to improve accuracy and robustness [12, 15].

The convergence of these two lines of research—federated learning and multi-generator spatiotemporal modeling—offers a compelling path toward privacy-preserving autonomous navigation prediction. By distributing the modeling task across multiple generators that are trained in a federated manner, the system can leverage the strengths of ensemble diversity while preserving data locality. The multi-generator framework can be designed to produce synthetic trajectory samples or predictive distributions that are shared across nodes, thereby enabling collaborative learning without direct exposure of private trajectories [10, 14]. This paper provides a comprehensive investigation of the architectural, governance, infrastructure, and policy dimensions of such a federated multi-generator spatiotemporal modeling system. We emphasize system-level trade-offs and practical deployment considerations rather than focusing solely on algorithmic improvements. The analysis draws on recent advances in federated optimization, differential privacy, secure aggregation, and trajectory prediction to construct a holistic view of the challenges and opportunities in this emerging domain.

2. Background and Related Work

The field of trajectory prediction has matured considerably over the past decade, with spatiotemporal graph-based models demonstrating strong performance in capturing both spatial interactions and temporal dependencies [2, 6]. Graph neural networks encode the relationships among agents at each time step, while recurrent or temporal convolutional layers model motion history. More recently, multi-generator architectures have been introduced to improve prediction diversity and handle multimodal futures [12, 15]. These models employ multiple generators, each specializing in a different motion pattern or behavioral cluster, and fuse their outputs through an aggregation or selection mechanism. The flexible multi-generator model proposed by Zhu et al. integrates fused spatiotemporal graph representations

to achieve robust trajectory predictions in complex urban environments [12]. Similarly, other works have explored adversarial training and variational inference to generate multiple plausible future trajectories [9, 11].

Parallel to these advances, federated learning has been extensively studied for applications in mobile sensing, healthcare, and autonomous driving [1, 4, 7]. Early federated learning algorithms, such as FedAvg, demonstrated that model accuracy could be maintained while reducing communication overhead and preserving data locality [1]. Subsequent research has focused on improving convergence rates in non-IID data distributions, handling client heterogeneity, and incorporating privacy guarantees through differential privacy [5, 8]. Secure aggregation protocols have been developed to ensure that the server cannot infer individual client updates even when aggregating encrypted parameters [3]. In the context of autonomous navigation, federated learning has been applied to object detection, path planning, and trajectory prediction, though most studies assume a single model architecture rather than a multi-generator ensemble [7, 13].

The intersection of federated learning and multi-generator modeling remains relatively underexplored. Preliminary works have proposed federated ensemble methods where each client maintains a collection of models and shares aggregated predictions [10, 14]. However, these approaches often require extensive communication of model parameters or synthetic data, which can strain bandwidth and raise additional privacy concerns. Our work builds on these foundations by proposing a systematic framework where multiple generators are trained collaboratively under strict privacy constraints, and their outputs are fused in a privacy-preserving manner. The design must balance the benefits of ensemble diversity against the increased communication and computation costs inherent in federated settings.

3. Architectural Framework and Design Trade-offs

The proposed federated multi-generator spatiotemporal modeling architecture consists of several key components: local nodes, each equipped with a multi-generator module and a local spatiotemporal encoder; a central aggregation server that coordinates model updates; and a privacy budget manager that enforces differential privacy constraints across training rounds. Each local node processes its own private trajectory data through a spatiotemporal graph network to extract latent representations, which are then fed into multiple generators. The generators produce either probabilistic trajectory distributions or a set of candidate trajectories. A fusion mechanism, such as a mixture density network or a learned attention weighting, combines the generator outputs into a final prediction.

The training process follows a federated optimization loop. In each round, the central server selects a subset of available nodes and sends the current global model parameters. Each selected node performs several local gradient descent steps on its private data using a loss function that incorporates both prediction accuracy and a penalty for deviations from the global model. After local training, the node sends back model updates, which are aggregated using a secure aggregation protocol to prevent the server from inspecting individual updates [3]. To enforce differential privacy, the server adds calibrated noise to the aggregated update before distributing it to clients in the next round [5]. The privacy budget is tracked globally and locally to ensure that the cumulative privacy loss remains below a predefined threshold.

Several design trade-offs emerge from this architecture. First, the number of generators per node introduces a trade-off between model capacity and communication cost. A larger number of generators increases the expressiveness of the local model but also multiplies the

number of parameters that must be transmitted and aggregated. One approach is to share only the parameters of the spatiotemporal encoder and the fusion mechanism while keeping the generator-specific parameters local [14]. However, this reduces the opportunity for cross-node learning of diverse motion patterns. Alternatively, nodes can share synthetic trajectory samples generated by their local ensemble, which may be less privacy-sensitive than raw data but still carry information leakage risks [10]. Second, the choice of fusion strategy affects both prediction quality and the vulnerability to adversarial manipulations. A weighted average of generator outputs is simple but may be dominated by a single generator that memorizes local patterns. More sophisticated fusion methods based on uncertainty estimation or out-of-distribution detection can improve robustness but require additional computation and training.

Another critical trade-off involves the frequency and granularity of communication. In federated learning, the convergence rate is influenced by the number of local epochs and the fraction of participating nodes. For autonomous navigation, latency constraints may require frequent model updates to adapt to rapidly changing environments, yet high communication frequency increases bandwidth consumption and the risk of privacy leakage through repeated updates. Adaptive communication strategies that adjust the number of local steps based on data heterogeneity or model uncertainty can help balance these considerations [7]. Additionally, the use of compressed model updates, such as gradient quantization or pruning, can reduce communication overhead without significantly degrading accuracy [4, 8].

4. Privacy-Preserving Mechanisms and Governance

Preserving the privacy of trajectory data in a federated multi-generator system requires a multi-layered governance framework that encompasses technical mechanisms, operational policies, and legal compliance. At the technical level, differential privacy is the most widely adopted approach for bounding the information leakage from model updates [5]. In our architecture, differential privacy is applied both locally during client-side training, where noise is added to compute gradients, and globally after aggregation. The privacy budget must be allocated across training rounds and across multiple generators, as each generator may have a different sensitivity to perturbations in the training data. Managing this allocation is particularly challenging when generators are specialized to rare or high-risk behaviors, such as emergency maneuvers, because a small privacy budget may cause the model to lose the ability to predict such events accurately.

Secure aggregation protocols provide an additional layer of protection by ensuring that the central server only observes the sum of client updates, not the individual contributions [3]. This prevents the server from performing membership inference or gradient inversion attacks based on a single client's update. However, secure aggregation requires clients to encrypt their updates before transmission, which introduces computational overhead and increases latency. In automotive contexts where vehicles may have limited computational resources, the practicality of full secure aggregation must be weighed against its privacy benefits. Alternative approaches, such as shuffling or random noise injection at the client level, can offer weaker but more efficient privacy guarantees [8].

Governance of privacy in a distributed system extends beyond algorithm design to include data use agreements, auditability, and accountability. Nodes participating in the federation must adhere to a common privacy policy that specifies the purpose of data processing, the retention period of model updates, and the rights of data subjects. Since trajectory data may include personally identifiable information such as license plates or location histories, the federation must implement data minimization principles, ensuring that only necessary

information is shared. One approach is to strip identifying metadata from trajectory sequences before they are used for local training, though this may reduce the richness of the representation. Another is to use synthetic trajectory generation as a privacy-enhancing technique, where nodes generate and share artificial trajectories that statistically resemble the real data without exposing exact paths [10, 14].

Audit trails are essential for verifying compliance with privacy regulations and for detecting potential breaches. In a federated system, each node can maintain a local log of the updates it has contributed, along with the corresponding differential privacy parameters. A global audit system can periodically verify that the aggregated model does not exhibit anomalous memory of specific trajectories. However, auditing itself raises privacy concerns, as the audit process may require inspection of model internals. Cryptographic techniques such as zero-knowledge proofs can enable audits without revealing sensitive information, but their deployment in resource-constrained environments remains an open challenge.

5. Infrastructure and Deployment Considerations

The success of a federated multi-generator system for autonomous navigation depends heavily on the underlying infrastructure, including communication networks, computational resources, and edge-cloud coordination. Autonomous vehicles generate massive amounts of trajectory data, often at high frequencies (10-30 Hz). Processing this data locally for model training requires onboard hardware capable of running deep neural networks with multiple generators. While modern vehicles are increasingly equipped with powerful GPUs and specialized AI accelerators, the computational budget for federated training must not interfere with real-time inference for navigation. Therefore, local training is typically performed during idle periods, such as when the vehicle is parked or charging, or on dedicated edge servers installed in smart infrastructure.

Communication infrastructure must support the periodic transmission of model updates between vehicles and a central aggregation server. In urban environments, 5G networks offer low latency and high bandwidth, enabling frequent communication rounds. However, in rural or underserved areas, connectivity may be intermittent and slow. The federated learning algorithm must be resilient to dropped connections and variable participation rates. Techniques such as asynchronous aggregation or client selection based on connectivity can mitigate these issues [7]. Additionally, the size of model updates can be reduced through compression, as mentioned earlier, but compression may introduce noise that interacts with the differential privacy mechanisms in unpredictable ways.

Edge computing nodes, such as roadside units or mobile network base stations, can serve as intermediate aggregators to reduce the load on the central server and improve latency. A hierarchical federation architecture, where vehicles first aggregate updates within a geographical region before sending a summary to a global server, can better scale to large fleets. This hierarchical approach also aligns with data sovereignty requirements, as regional aggregators can enforce local privacy regulations before forwarding data to central servers located in different jurisdictions.

Sustainability is another critical infrastructure dimension. The energy consumption of training deep models on millions of vehicles, even if only periodically, can be substantial. Federated learning reduces the need for data transmission to a centralized data center, but the energy cost of local computation and secure aggregation must be considered. Carbon-aware scheduling, where training is performed during periods of low grid carbon intensity or when

renewable energy is abundant, can mitigate environmental impact [16]. Furthermore, the hardware lifecycle of vehicles and edge devices must be planned to avoid e-waste, especially as models evolve and require hardware upgrades.

6. Robustness, Fairness, and Sustainability

Robustness in federated multi-generator systems is threatened by several factors: non-IID data distributions across nodes, adversarial attacks on model updates, and natural distributional shifts due to changes in traffic patterns or weather conditions. Multi-generator architectures inherently provide robustness through diversity, as different generators may capture complementary patterns that compensate for missing or noisy data. However, if the federated training process biases the generators toward a common mode (e.g., because of over-regularization or strong model averaging), diversity may be lost. Techniques such as personalized federated learning, where each node tailors its local model to its own data distribution while still benefiting from global knowledge, can preserve diversity [17].

Adversarial attacks on federated learning, such as Byzantine attacks where malicious nodes submit corrupted updates, can compromise the global model. Multi-generator systems can mitigate this by employing robust aggregation rules, such as median-based or trimmed means, that are less sensitive to outliers [18]. Additionally, the fusion mechanism can be designed to down-weight generators that produce implausible predictions, acting as a defense against adversarial inputs.

Fairness is a multidimensional concept in autonomous navigation. The predictions generated by the federated model must be accurate across different geographic regions, demographic groups, and types of road users (e.g., pedestrians, cyclists). If the federation is dominated by nodes from affluent neighborhoods with modern infrastructure, the model may perform poorly in underserved areas with different traffic patterns. This raises concerns about algorithmic bias and equitable access to the benefits of autonomous navigation technology. To address fairness, the aggregation process can incorporate importance weights that give higher representation to underrepresented nodes, or the federation can be required to include a diverse set of nodes from various socioeconomic contexts [19]. Additionally, the privacy guarantees themselves may need to be adaptive: nodes with more sensitive data may require stronger privacy protections, but this could lead to disparity in model accuracy if those nodes contribute noisier updates.

Sustainability, as introduced earlier, intersects with robustness and fairness. For example, a system that prioritizes energy efficiency by reducing the number of training rounds may achieve lower accuracy for certain nodes, exacerbating fairness issues. Conversely, a system that requires all nodes to train equally may incur high energy costs and accelerate hardware depreciation. A multi-objective optimization framework that balances prediction accuracy, privacy loss, energy consumption, and fairness metrics is needed for responsible deployment.

7. Case Illustrations and Cross-Domain Comparisons

To contextualize the federated multi-generator approach, we compare it with three alternative paradigms: centralized spatiotemporal modeling, federated single-generator modeling, and local-only modeling (no federation). In centralized modeling, all trajectory data is aggregated at a central server, where a multi-generator model is trained with full access to the global data distribution. This approach typically achieves the highest prediction accuracy because it can capture rare or region-specific patterns without noise from privacy mechanisms. However, it

suffers from severe privacy risks and regulatory non-compliance, as raw data is exposed. Moreover, the central server becomes a single point of attack and failure.

Federated single-generator modeling, where each node trains the same architecture (e.g., a single spatiotemporal graph network) and shares updates, reduces privacy risk but may lose the diversity benefits of multiple generators. In tasks with multimodal trajectory distributions, a single generator often produces overconfident or unimodal predictions, leading to higher error rates in intersection and roundabout scenarios [9]. The multi-generator framework addresses this limitation by maintaining multiple hypotheses, which can be combined to produce a richer predictive distribution.

Local-only modeling, where each node trains a model entirely on its own data without federation, preserves maximum privacy but suffers from poor data efficiency and inability to generalize to unseen scenarios. In the event of a rare event (e.g., a pedestrian crossing an unmarked street), a local model may have no training examples and fail to predict it. The federated multi-generator architecture offers a middle ground, enabling knowledge transfer across nodes while keeping raw data private. The trade-off is that the prediction accuracy may be slightly lower than the centralized baseline due to privacy noise and communication inefficiencies.

A case study of an urban region with mixed traffic (cars, buses, bicycles) illustrates these trade-offs. In such a setting, the multi-generator model can allocate one generator to handle bus trajectories that follow fixed routes, another for bicycle movements that are less predictable, and a third for general car traffic. Under federation, nodes from bus-heavy districts can contribute knowledge about bus patterns, while nodes from bike-sharing hubs contribute bicycle patterns, without revealing the exact GPS traces of individuals. The fusion mechanism must decide how to weight each generator globally, but regional variations may require personalized fusion weights.

8. Future Directions and Policy Implications

The deployment of federated multi-generator spatiotemporal modeling for autonomous navigation prediction is still in its early stages, and several open research and policy challenges remain. From a technical perspective, developing efficient methods for training multiple generators under heterogeneous privacy budgets is a priority. Each generator may need to be trained with different noise levels depending on the sensitivity of the behaviors it models. Adaptive privacy budget allocation that learns which generators are most informative without overburdening the privacy budget is an exciting direction.

Another promising avenue is the integration of synthetic data sharing alongside model updates. Nodes could generate synthetic trajectories using their local multi-generator models and share these with other nodes or a central server, allowing for data augmentation without exposing real trajectories. The synthetic data must satisfy differential privacy guarantees itself, which can be achieved through techniques such as differentially private generative adversarial networks [20]. The combination of federated multi-generator learning with synthetic data generation could enable more efficient knowledge transfer while maintaining strong privacy protections.

Policy implications are profound. Autonomous navigation systems are increasingly regulated, with governments mandating safety standards and data protection measures. The federated multi-generator framework could serve as a reference architecture for compliance, as it inherently supports data minimization and purpose limitation. However, regulators must also

address the challenge of auditing such distributed systems. Standardization bodies, such as the IEEE and ISO, are beginning to develop guidelines for federated learning in autonomous systems, but concrete norms for multi-generator models are lacking. Policies should incentivize the adoption of privacy-preserving architectures through certification schemes and liability frameworks that differentiate between centralized and decentralized approaches.

Finally, ethical considerations regarding the potential for surveillance or discrimination must be proactively addressed. Even in federated settings, the aggregated model might inadvertently encode biases from certain nodes if not carefully governed. Transparent reporting of the composition of the federation, the privacy budgets used, and the performance across different demographic and geographic groups should become standard practice. The research community, industry consortia, and policymakers must collaborate to ensure that federated multi-generator systems for autonomous navigation are not only technically effective but also socially just and environmentally sustainable.

9. Conclusion

This paper has presented a comprehensive analysis of federated multi-generator spatiotemporal modeling for privacy-preserving autonomous navigation prediction. By integrating multiple generative trajectory predictors within a federated learning framework, the proposed architecture addresses the critical need for high-accuracy trajectory forecasting while respecting data privacy and regulatory constraints. We examined structural trade-offs in the number of generators, fusion strategies, and communication frequency, as well as governance mechanisms including differential privacy, secure aggregation, and auditability. Infrastructure considerations such as edge computing, network heterogeneity, and energy sustainability were discussed alongside robustness, fairness, and policy implications. Comparative case illustrations highlighted the advantages of the federated multi-generator approach over centralized, single-generator, and local-only alternatives. As autonomous navigation systems become ubiquitous, the federated multi-generator paradigm offers a viable path toward a future where advanced prediction capabilities do not come at the cost of individual privacy or societal equity. Continued research will refine the architecture, validate its performance in real-world deployments, and inform the development of ethical and regulatory frameworks.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (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.

2. Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In Proceedings of the 27th International Joint Conference on Artificial Intelligence (pp. 3634-3640).

3. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & 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).

4. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.

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).

6. Li, J., Yang, F., Tomizuka, M., & Choi, C. (2020). EvolveGraph: Multi-agent trajectory prediction with dynamic relational reasoning. In Advances in Neural Information Processing Systems (Vol. 33, pp. 19783-19794).

7. Wang, J., Charles, Z., Xu, Z., Joshi, G., McMahan, H. B., & al-Kashif, M. (2019). A field guide to federated optimization. arXiv preprint arXiv:1910.03218.

8. Papernot, N., Abadi, M., Erlingsson, Ú., Goodfellow, I., & Talwar, K. (2017). Semi-supervised knowledge transfer for deep learning from private training data. arXiv preprint arXiv:1610.05755.

9. Rhinehart, N., McAllister, R., & Levine, S. (2019). Deep imitative models for flexible inference, planning, and control. arXiv preprint arXiv:1904.00999.

10. Augenstein, S., McMahan, H. B., Ramage, D., Ramaswamy, S., & Thakurta, A. (2019). Generative models for effective ML on private, decentralized datasets. arXiv preprint arXiv:1904.09620.

11. Schmerling, E., Leung, K., & Pavone, M. (2022). Multimodal probabilistic trajectory prediction via adversarially learned generative models. In Proceedings of the 2022 International Conference on Robotics and Automation (pp. 11096-11102).

12. Zhu, P., Han, F., & Deng, H. (2023, December). Flexible multi-generator model with fused spatiotemporal graph for trajectory prediction. In IET Conference Proceedings CP874 (Vol. 2023, No. 47, pp. 417-422). Stevenage, UK: The Institution of Engineering and Technology.

13. Liang, L., Chen, Y., & Liao, Q. (2021). Federated learning for autonomous driving: A survey. IEEE Transactions on Intelligent Vehicles, 6(4), 712-726.

14. Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Fan, F., & Shen, Z. (2020). Federated learning with differential privacy: Algorithms and performance analysis. IEEE Transactions on Information Forensics and Security, 15, 3454-3469.

15. Chai, Y., Zhu, P., & Deng, H. (2023). Multi-generator trajectory prediction with adaptive fusion. IEEE Transactions on Intelligent Transportation Systems, 24(8), 8742-8753.

16. Qiu, J., Wu, Q., Ding, G., Xu, Y., & Feng, S. (2016). A survey of machine learning for big data processing. EURASIP Journal on Advances in Signal Processing, 2016(1), 67.

17. 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.

18. Blanchard, P., Guerraoui, R., Stainer, J., & others. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. In Advances in Neural Information Processing Systems (Vol. 30).

19. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. In Advances in Neural Information Processing Systems (Vol. 29).

20. Xie, L., Lin, K., Wang, S., Wang, F., & Zhou, J. (2018). Differentially private generative adversarial network. arXiv preprint arXiv:1802.06739.