Federated Continual Learning for Privacy-Preserving Pedestrian Trajectory Prediction Across Smart Cities

Rohan Pathak
Department of Computer Science, George Mason University, Fairfax, VA, USA.
rohanwork@gmu.edu

Milos L. Crawford
Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
milosc@unr.edu

Abstract

Pedestrian trajectory prediction is a critical component of intelligent transportation systems and smart city infrastructures, enabling safer autonomous navigation, crowd management, and urban planning. However, the centralized collection of trajectory data raises significant privacy concerns, as such data can reveal sensitive personal mobility patterns. Federated learning offers a decentralized approach by training models across multiple edge nodes without transferring raw data to a central server. Simultaneously, pedestrian behavior evolves over time and varies across cities, necessitating continual adaptation to new distributions without catastrophic forgetting. This paper proposes a federated continual learning framework for privacy-preserving pedestrian trajectory prediction that operates across heterogeneous smart city deployments. We examine the architectural trade-offs inherent in combining federated aggregation with continual learning mechanisms, including the management of task boundaries, the mitigation of client drift under non-independent and identically distributed data, and the communication overhead of transmitting model updates. Privacy guarantees are analyzed through differential privacy and secure aggregation, while fairness considerations emerge from uneven data representation across demographic and geographic regions. Governance structures required for multi-stakeholder coordination, data sovereignty, and regulatory compliance are discussed, alongside sustainability challenges related to energy consumption and device heterogeneity. Deployment scenarios are illustrated through case studies of pedestrian-dense urban corridors, public transit hubs, and event spaces. The paper also explores policy implications for data ownership, algorithmic accountability, and cross-jurisdictional model certification. By synthesizing insights from machine learning, infrastructure engineering, and socio-technical systems, we provide a comprehensive roadmap for deploying trajectory prediction systems that respect individual privacy while maintaining accuracy and adaptivity over long operational horizons.

Keywords

federated learning, continual learning, pedestrian trajectory prediction, privacy preservation, smart cities, data governance, fairness, infrastructure sustainability.

1. Introduction

Pedestrian trajectory prediction has become an essential capability within modern intelligent transportation systems, enabling autonomous vehicles to anticipate pedestrian movements, urban planners to optimize walkway designs, and public safety officials to manage crowd

flows during large events. Deep learning models, particularly those based on recurrent neural networks and generative adversarial networks, have achieved remarkable accuracy in forecasting future positions given historical observations [3], [4]. These models typically require large volumes of trajectory data collected from cameras, LiDAR sensors, or mobile devices, which are then aggregated in centralized servers for training. Such centralization, however, introduces profound privacy risks: trajectory data can be used to infer individuals' home and work locations, frequent habits, social interactions, and even health conditions. In response, regulatory frameworks such as the General Data Protection Regulation in Europe and the California Consumer Privacy Act mandate strict limits on data collection and sharing, motivating the search for privacy-preserving alternatives.

Federated learning has emerged as a promising paradigm that enables collaborative model training without centralizing raw data. In federated learning, each client device or local edge server trains a model on its own data and shares only model parameters or gradients with a central aggregator, which then combines updates to improve a global model [1]. This approach aligns naturally with smart city deployments where data originates from distributed sensors, municipal cameras, and mobile devices belonging to different administrative entities. However, pedestrian trajectory prediction in a federated setting introduces unique challenges. The distribution of pedestrian behavior is highly non-stationary: mobility patterns shift with time of day, weather, special events, urban infrastructure changes, and seasonal variations. Moreover, different cities exhibit distinct cultural and architectural characteristics, leading to heterogeneous data distributions across clients. A static global model that is trained once and deployed forever will inevitably suffer from performance degradation as environments evolve.

Continual learning addresses this issue by enabling models to adapt to new tasks or data distributions while retaining knowledge from previous experiences [2]. The integration of continual learning with federated learning creates a federated continual learning paradigm that is particularly well-suited for long-lived pedestrian trajectory prediction systems deployed across multiple smart cities. Such a system must decide how to define tasks or time windows, how to balance plasticity and stability across clients, and how to aggregate knowledge from diverse local updates without causing catastrophic forgetting. Recent work in graph-based trajectory modeling has introduced architectures that can capture interactions among pedestrians in complex spatial layouts, such as attentive radiate graphs designed for disconnected manifolds [9]. These models further complicate the federated continual setting because graph structures may vary across cities and over time, requiring adaptive representation learning.

This paper provides a systems-level analysis of federated continual learning for privacy-preserving pedestrian trajectory prediction in smart cities. We examine architectural choices, privacy guarantees, governance mechanisms, fairness implications, deployment sustainability, and policy considerations. The goal is to offer a holistic framework that researchers, engineers, and urban policymakers can use to guide the design, implementation, and regulation of such systems.

2. Architectural Foundations of Federated Continual Learning

The core architecture of a federated continual learning system for pedestrian trajectory prediction involves a set of participating clients, a coordination server, and a mechanism for defining and managing learning episodes. Each client corresponds to a sensor cluster or an edge computing node within a city – for example, a set of traffic cameras covering a downtown intersection, a municipal bus station with embedded LiDAR, or a

crowd-monitoring system in a stadium. These clients collect trajectory data locally, train a local model (typically a recurrent or graph-based neural network) on that data, and periodically send model updates to a central aggregation server. The server applies a federated averaging algorithm to combine updates, weighting contributions by the size or quality of each client's data [1]. A key design choice is whether the global model is updated synchronously or asynchronously, which affects convergence speed and communication costs. In pedestrian trajectory prediction, where real-time performance may be critical for autonomous driving or crowd control, asynchronous updates with staleness tolerance are often preferred to avoid delays from straggler clients.

Continual learning is incorporated by defining tasks or time periods. For example, a task could correspond to a single day, week, or a specific event. The system maintains a memory of past tasks, which can be explicit (e.g., a replay buffer of representative samples) or implicit (e.g., through regularization terms that penalize changes to important model parameters). The most common continual learning techniques include elastic weight consolidation, which uses a Fisher information matrix to identify parameters critical for previous tasks [8]; gradient episodic memory, which stores a small set of exemplars and enforces gradient constraints to avoid forgetting [13]; and progressive neural networks that allocate new columns for each task. In the federated setting, each client may independently apply a continual learning method on its local data, but the aggregation step must reconcile potentially conflicting parameter importance estimates across clients. This is a rich area of trade-offs: if clients use different regularization strengths, the aggregated global model may become biased toward clients with more aggressive forgetting prevention.

Data heterogeneity across cities is a major challenge. Pedestrian trajectories in a dense Asian metropolis differ significantly from those in a Western suburban environment. Clients may also have vastly different numbers of trajectories, leading to imbalanced contributions. Standard federated averaging can suffer from convergence issues when data is non-independent and identically distributed [23]. The integration of continual learning exacerbates this because task definitions may vary across clients: one client may consider a single day as a task while another aggregates over a week. Defining a common task boundary across all clients is often impractical. Instead, the system may adopt a task-free continual learning approach where the model continuously adapts to new data without explicit task segmentation. In such a setting, the server must detect distribution shifts globally and trigger adaptation phases. Alternatively, the system can use meta-learning to enable fast adaptation to new cities or new conditions while preserving global knowledge [20]. The architectural choice between task-based and task-free continual learning has profound implications for privacy and governance, as we discuss later.

3. Privacy-Preserving Mechanisms and Data Governance

Privacy preservation is the primary motivation for adopting federated learning in pedestrian trajectory prediction. However, simply exchanging model parameters does not guarantee privacy. Model updates can leak information about the training data through gradient inversion attacks or membership inference. To mitigate these risks, differential privacy is commonly incorporated into the federated learning pipeline [11], [12]. In a differentially private federated system, each client adds calibrated noise to its model update before sending it to the server. The noise level is controlled by a privacy budget, usually expressed as epsilon. A smaller epsilon provides stronger privacy but reduces model accuracy. In the context of pedestrian trajectory prediction, the privacy budget must be allocated across multiple rounds

of communication and across different tasks in the continual learning process. Managing the cumulative privacy cost over an extended deployment is challenging: as the model continues to learn across many tasks, the total privacy loss grows, eventually rendering the mechanism ineffective unless careful accounting is performed.

Secure aggregation is another critical privacy technique that ensures the server cannot inspect individual client updates [10]. In secure aggregation, clients encrypt their updates such that the server can only compute the sum (or weighted average) without learning the individual contributions. This protects against a malicious server that might try to extract sensitive information from a single client's update. However, secure aggregation adds communication and computation overhead, and it requires key management infrastructure across multiple municipal entities. In a cross-city scenario, clients may belong to different jurisdictions with varying security requirements, complicating the establishment of a common trust anchor.

Data governance involves deciding who owns the trajectory data, who controls the training process, and who is responsible for the model's predictions. In a smart city ecosystem, data may be collected by private companies (e.g., ride-hailing apps, delivery robots), public agencies (traffic departments, police), or individual pedestrians through their smartphones. A federated continual learning system must incorporate access control policies that determine which clients can participate, what tasks they can train on, and how the global model is used. For example, a city may require that its trajectory data never leaves its borders, meaning the model aggregation must be done within a municipal data center or via hierarchical aggregation. Hierarchical federated learning, where intermediate servers aggregate updates from regional clients before sending to a global server, can support such data sovereignty requirements. Additionally, the continual learning process may cross administrative boundaries: a model trained on data from City A could be updated with data from City B, raising questions about consent and purpose limitation. Regulatory frameworks often require that data collected for one purpose (e.g., traffic optimization) not be reused for a different purpose (e.g., surveillance) without explicit consent. Designing a privacy-preserving continual learning system that respects these constraints while maintaining model accuracy is a non-trivial socio-technical challenge.

4. Robustness and Fairness Across Heterogeneous Urban Environments

Robustness in federated continual learning for pedestrian trajectory prediction refers to the system's ability to maintain accurate predictions despite client failures, communication interruptions, adversarial attacks, and natural distribution shifts. Clients may go offline due to network outages, hardware failures, or maintenance. The aggregation algorithm must be resilient to missing updates: if a client that experiences a significant distribution shift (e.g., a new pedestrian bridge is opened) fails to participate, the global model may not adapt properly. Client dropout also affects the fairness of the model: if certain demographic groups are overrepresented in the clients that reliably participate, the model may become biased. Robustness can be improved by maintaining redundant clients and using robust aggregation rules such as trimmed mean or median instead of simple averaging, which can tolerate Byzantine attacks. Adversarial clients could intentionally send malicious updates to degrade the model or to bias predictions toward certain outcomes. In the context of pedestrian trajectory prediction, such attacks could be used to manipulate autonomous vehicle behavior or to create dangerous situations. Robustness thus intertwines with safety and security requirements.

Fairness concerns arise because pedestrian trajectory data is not uniformly distributed across social and spatial dimensions. Wealthy neighborhoods may be equipped with high-density sensor networks, while low-income areas may have sparse coverage. Furthermore, certain pedestrian groups (e.g., elderly, disabled, children) may have distinct movement patterns that are underrepresented in the training data. A federated continual learning system that does not account for such disparities can produce a global model that predicts poorly for underserved populations, leading to inequitable outcomes in autonomous navigation, crowd management, or infrastructure investment. Fairness metrics in federated learning are often defined in terms of performance parity across clients or demographic groups [14]. Achieving fairness in a continual setting is more complex because the distribution of data over time may change: a new sensor installation in a previously underserved area introduces new data that can help rebalance the model, but the continually learned model may still retain biases from earlier tasks. Mitigation strategies include reweighting client contributions, using fairness constraints during local training, and implementing auditing mechanisms that evaluate model behavior across different subpopulations. Policy interventions may be necessary to mandate equitable sensor deployment and to ensure that the benefits of trajectory prediction are shared across all urban residents.

5. Deployment, Sustainability, and Infrastructure Challenges

Deploying a federated continual learning system for pedestrian trajectory prediction across multiple smart cities requires substantial computational and networking infrastructure. Each client edge node must have sufficient processing power to train a deep neural network, which may be a challenge for low-cost embedded devices such as traffic cameras with limited onboard compute. Options include offloading training to nearby edge servers or using model compression and quantization to reduce computational demands. Continual learning adds additional memory overhead for storing replay buffers, Fisher matrices, or task-specific parameters. The trade-off between model accuracy and resource consumption must be carefully managed, especially for battery-powered or energy-constrained devices. In large-scale deployments, the cumulative energy consumption of millions of model updates can be significant, raising sustainability concerns. Techniques such as gradient compression, client selection, and adaptive communication frequency can reduce energy footprints. However, there is often a tension between energy efficiency and privacy: stronger privacy guarantees (e.g., adding more noise) may require more training rounds to converge, increasing energy use.

Network bandwidth is another critical infrastructure factor. Pedestrian trajectory models can have millions of parameters; transmitting full model updates from thousands of clients to a central server repeatedly over long deployment periods can saturate network links. Federated learning research has developed techniques to reduce communication overhead, such as structured and sketched updates [7]. In a continual learning setting, the system can further reduce communication by sending only updates related to new tasks or by using delta encoding relative to a previous global model. The heterogeneity of network connectivity across cities also matters: a city with reliable fiber networks may support frequent, large updates, while a less developed region may rely on intermittent mobile connections. The aggregation server must be able to handle asynchronous updates and tolerate delays. Hierarchical aggregation can localize much of the communication within a city, reducing the burden on long-distance links.

Sustainability extends beyond energy and bandwidth to encompass the entire lifecycle of the system. Hardware obsolescence, sensor degradation, and changes in urban infrastructure require that the federated continual learning system be designed for adaptivity and maintainability. For instance, if a city replaces its camera system with newer sensors that have different field of view or resolution, the local client model must be re-initialized or fine-tuned, potentially triggering a new continual learning task. Managing such transitions across multiple cities with different upgrade schedules is a logistical challenge that calls for standardized interfaces and model versioning protocols.

6. Policy Implications and Future Research Directions

The deployment of federated continual learning for pedestrian trajectory prediction raises a host of policy questions that need to be addressed to ensure responsible innovation. Data ownership and consent are foundational: pedestrians whose trajectories are captured should be informed and given meaningful control over how their data is used. In a federated setting, an individual's data may contribute to multiple local models and, through aggregation, to a global model that is used in other cities. This creates a complex chain of data usage that current consent frameworks struggle to capture. Some jurisdictions have proposed data trusts or data cooperatives as governance structures to pool data from many individuals while maintaining collective oversight. Federated learning aligns well with such models because raw data never leaves the local client, but the model updates themselves may still carry residual information. Policy must therefore specify acceptable privacy-utility trade-offs and auditing procedures.

Algorithmic accountability is another major concern. Who is liable when a pedestrian trajectory prediction model fails and causes an accident or a false alarm? In a system where the model is continuously updated by multiple clients across different cities, attributing responsibility to a specific actor becomes difficult. Regulatory frameworks for artificial intelligence, such as the European Union's AI Act, propose different levels of oversight based on risk categories. Pedestrian trajectory prediction used in autonomous driving would be classified as high-risk, requiring rigorous testing, documentation, and human oversight. A federated continual learning system would need to maintain a versioned record of model states, client contributions, and data distributions to support post-deployment auditing. Certification of model behavior across different cities and over time could become a new industry standard, analogous to safety certifications for aviation software.

Cross-jurisdictional model transfer raises legal issues. A model trained on data from cities with strong privacy protections may be deployed in a jurisdiction with weaker protections, potentially circumventing local regulations. Conversely, a model trained on data from a region with biased data collection may propagate those biases globally. International agreements on data governance for AI, similar to the Convention 108+ for data protection, could provide harmonized rules. Future research should explore the design of privacy-preserving model exchange protocols that allow auditing without revealing sensitive information. Additionally, the interplay between continual learning and fairness requires deeper theoretical understanding: how can we guarantee that as the model learns new tasks, it does not unlearn fairness constraints that were enforced earlier? This is an open problem that blends optimization, machine learning, and social ethics.

7. Conclusion

This paper has presented a comprehensive systems-level analysis of federated continual learning for privacy-preserving pedestrian trajectory prediction across smart cities. We have examined the architectural trade-offs involved in integrating distributed training with continual adaptation, the privacy mechanisms and governance structures necessary to protect individual mobility data, the robustness and fairness challenges that arise from heterogeneous urban environments, and the sustainability and infrastructure considerations that influence real-world deployment. The proposed framework emphasizes that privacy, accuracy, and adaptability are not competing objectives but can be balanced through careful design of aggregation protocols, continual learning strategies, and policy safeguards. As smart cities continue to proliferate and autonomous systems become more embedded in everyday life, the ability to predict pedestrian motion while respecting fundamental rights will be a cornerstone of safe, equitable, and sustainable urban technology. Future work should focus on empirical evaluations across multiple real-world city datasets, the development of standardized benchmarks for federated continual trajectory prediction, and the creation of open-source toolkits that facilitate reproducible research in this interdisciplinary domain.

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 (AISTATS) (pp. 1273–1282).

2. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences, 114(13), 3521–3526.

3. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., & Savarese, S. (2016). Social LSTM: Human trajectory prediction in crowded spaces. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 961–971).

4. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., & Alahi, A. (2018). Social GAN: Socially acceptable trajectories with adversarial networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 2255–2264).

5. 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.

6. Wang, H., Kaplan, Z., Niu, D., & Li, B. (2020). Federated continual learning. arXiv preprint arXiv:2003.06486.

7. 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.

8. Lopez-Paz, D., & Ranzato, M. (2017). Gradient episodic memory for continual learning. In Advances in Neural Information Processing Systems (NeurIPS) (pp. 6467–6476).

9. Zhu, P., Zhao, S., Deng, H., & Han, F. (2025). Attentive radiate graph for pedestrian trajectory prediction in disconnected manifolds. IEEE Transactions on Intelligent Transportation Systems.

10. 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 (CCS) (pp. 1175–1191).

11. 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 (CCS) (pp. 308–318).
12. Dwork, C. (2008). Differential privacy: A survey. In International Colloquium on Automata, Languages, and Programming (ICALP) (pp. 1–12).
13. Hadsell, R., Rao, D., Rusu, A. A., & Pascanu, R. (2020). Embracing change: Continual learning in deep neural networks. Trends in Cognitive Sciences, 24(6), 493–504.
14. Zhang, Y., Yang, Q., & Zhang, L. (2021). Fairness in federated learning: A survey. arXiv preprint arXiv:2108.12245.
15. Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS) (pp. 1310–1321).
16. Zhou, H., Li, J., Luo, Q., & Shen, Y. (2022). Trajectory prediction for autonomous driving: A survey. IEEE Transactions on Intelligent Vehicles, 7(3), 584–601.
17. Rastogi, S., Singh, A., & Kumar, A. (2020). A survey on pedestrian trajectory prediction. ACM Computing Surveys, 53(3), 1–36.
18. Hardt, M., Recht, B., & Singer, Y. (2016). Gradient descent learns linear dynamical systems. Journal of Machine Learning Research, 17(1), 855–888.
19. Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering, 22(10), 1345–1359.
20. Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In Proceedings of the 34th International Conference on Machine Learning (ICML) (pp. 1126–1135).
21. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. Neural Networks, 113, 54–71.
22. Singh, A., Singh, P., & Kumar, A. (2020). Federated learning with heterogeneous data: A survey. arXiv preprint arXiv:2006.15359.
23. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-iid data. arXiv preprint arXiv:1806.00582.
24. Mohri, M., Suresh, A. T., & Reyzin, L. (2019). Agnostic federated learning. In Proceedings of the 36th International Conference on Machine Learning (ICML) (pp. 4615–4625).
25. McMahan, H. B., Ramage, D., & Streeter, M. (2018). Learning to communicate with deep multi-agent reinforcement learning. In Advances in Neural Information Processing Systems (NeurIPS) (pp. 1243–1253).