

Deep Learning-Based Spatiotemporal Modeling of Bicycle Crash Risk Using Mobile Sensing and Urban Built Environment Data

Yuzhuang Tang

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

yu.tang698@oregonstate.edu

Rishi L. Nair

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
rishi1991@buffalo.edu

Traver Schmidt

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence,
KS, USA.

schmidttravor@ku.edu

Abstract

Bicycle transportation is increasingly recognized as a sustainable urban mobility mode, yet pervasive safety concerns inhibit realization of its full societal potential. Conventional crash analysis methods rest on aggregate police-reported datasets that suffer from underreporting, spatiotemporal sparseness, and limited characterization of near-miss events. This paper presents a system-level investigation of deep learning architectures designed for spatiotemporal bicycle crash risk modeling by integrating heterogeneous streams of mobile sensing data and detailed urban built environment information. The discussion is oriented around structural trade-offs inherent in sensor fusion, computational graph construction, attention-driven memory mechanisms, and adversarial robustness. Particular emphasis is placed on the governance of large-scale mobility data pipelines, including privacy preservation, algorithmic fairness, infrastructure interoperability, and long-term deployment sustainability. Drawing from cross-domain comparisons with motor vehicle crash prediction and urban flow forecasting, we analyze the interplay between model complexity, interpretability, and real-world policy actionability. We further examine how varying spatial resolutions, temporal update frequencies, and built environment feature encodings influence distributive equity in risk assessments across sociodemographic groups. The paper concludes with an exploration of anticipatory governance frameworks that reconcile the promise of proactive infrastructure safety interventions with the public accountability demands of municipal planning systems, offering a forward-looking perspective on deep learning as an infrastructural component for sustainable urban safety systems.

Keywords

bicycle safety, spatiotemporal deep learning, mobile sensing, built environment, urban informatics, algorithmic fairness, system governance.

1. Introduction

The promotion of cycling as a mainstream mode of urban transport is widely endorsed for its capacity to reduce traffic congestion, lower greenhouse gas emissions, and improve public health outcomes [1]. However, the stark disparity in injury and fatality rates between cyclists and occupants of motorized vehicles continues to pose a formidable barrier to widespread adoption [2]. While extensive research has identified road geometry, land use patterns, traffic volumes, and visibility conditions as salient determinants of bicycle crash occurrence, the inherently dynamic and context-sensitive nature of cycling risk has remained difficult to capture through traditional statistical modeling approaches. Conventional safety analyses typically rely on historical crash records that are temporally aggregated and spatially coarse, often missing the subtle interactions between moving cyclists, momentary environmental conditions, and infrastructure design features. The result is a reactive planning paradigm in which infrastructure improvements follow fatal collisions rather than preemptively mitigating latent risk.

Recent advances in mobile sensing, urban computing, and deep learning present an unprecedented opportunity to operationalize highly granular, spatiotemporally resolved bicycle crash risk models that can learn from large-scale, heterogeneous data streams. Mobile phones, wearable devices, and dedicated cycling sensors can capture continuous kinematic traces, location fixes, and physiological signals, while municipal open data platforms increasingly provide rich representations of the built environment, including street network topology, intersection control types, bicycle facility types, land use mix, and micro-scale design features [3]. The convergence of these data modalities invites the design of deep neural architectures that model the spatial dependencies, temporal dynamics, and contextual conditioning of risk in ways that traditional statistical models cannot. Yet, the transition from proof-of-concept models to operationally deployed systems demands careful consideration of a range of structural trade-offs that transcend pure predictive accuracy. This paper provides a system-level analysis of deep learning-based spatiotemporal modeling of bicycle crash risk, emphasizing the architectural, governance, infrastructural, and equity dimensions that must be navigated for these techniques to become trusted components of municipal safety management frameworks.

2. Background and Related Work

Cycling safety research has historically drawn from the domain of road traffic crash analysis, which has developed a rich methodological toolkit spanning Poisson-gamma count models, logit-based severity models, and spatially correlated random parameter frameworks. However, the transferability of motor vehicle-centric models to bicycle safety contexts is limited by fundamental differences in exposure measurement, vulnerability, and crash reporting biases. The built environment plays a distinct role for cyclists: factors such as network connectivity, street lighting, separation from motorized traffic, and intersection density exert influences that are often non-linear and highly localized. Meta-analytic syntheses have confirmed that elasticities between travel behavior and built environment variables vary substantially by mode and geography, complicating any direct extrapolation of car-oriented safety findings to bicycling [3]. Moreover, the dependent variable itself is fraught with measurement challenges. Police-reported bicycle crashes represent a small and biased fraction of total safety-critical events, with underreporting strongly correlated with cyclist demographics, injury severity, and jurisdictional reporting practices [2]. The concept of a “safety continuum” that includes near-misses and evasive maneuvers has therefore gained traction, enabled by novel data collection methodologies.

The introduction of naturalistic cycling studies, where volunteer participants ride instrumented bicycles and generate continuous kinematics and video data, marked a significant shift toward proactive risk measurement [11]. These datasets allow the identification of safety-critical events from braking patterns, swerving, and proximity conflicts that never appear in official records. Simultaneously, the proliferation of smartphones equipped with Global Positioning System receivers, accelerometers, gyroscopes, and barometers has democratized the collection of mobility traces at population scale. Mobile phone data, when aggregated appropriately, have been used to infer travel modes, route choice, and urban activity patterns with high temporal resolution [17]. The challenge, however, lies in moving beyond descriptive analytics to predictive models that can anticipate where and when bicycle crashes or near-misses are most likely to occur, thereby enabling just-in-time interventions ranging from dynamic signage to adaptive traffic signal timing.

Deep learning has reshaped the landscape of spatiotemporal prediction across numerous urban domains. Foundational architectures such as convolutional neural networks, recurrent neural networks, and their hybrid variants have demonstrated superiority in capturing complex non-linear dependencies in tasks like traffic flow forecasting and crowd density prediction [5]. The explicit incorporation of spatial structure through graph convolutional networks enabled learning on irregular topologies such as road networks, where the adjacency relationships among segments encode critical connectivity information [6]. Attention mechanisms, originally developed for natural language processing, have been repurposed for spatiotemporal problems to capture long-range dependencies and dynamically weight the relevance of past observations and neighboring locations [7]. These developments have spurred a wave of applications in traffic crash prediction, particularly on expressways where loop detector and radar data provide dense sensor coverage [9]. The adaptation of such architectures to bicycle crash risk prediction, however, remains nascent. The irregular, low-density, and multi-modal nature of cycling data introduces distinctive challenges related to spatial autocorrelation scales, temporal intermittency, and the fusion of built environment features that vary in granularity from street-level design attributes to neighborhood-level census variables.

Furthermore, methodological frontier analyses in accident research have underscored the need for modeling frameworks that accommodate unobserved heterogeneity, non-stationary processes, and complex substitution patterns among risk factors [10]. The deep learning paradigm, with its capacity for hierarchical representation learning, holds promise for automatically disentangling these interactions, but its application must be carefully aligned with the system context: data quality, sensor power constraints, latency requirements, and the institutional capacity to act on risk score outputs. The literature contains numerous examples of successful deep learning systems for vehicular crash prediction on high-speed roads [9], and the extension to bicycle safety requires grappling with the finer spatiotemporal granularity and richer environmental context that characterizes urban cycling. Recent work has begun to harness emerging mobile sensing data specifically for bicycle safety risk assessment, demonstrating the feasibility of integrating crowdsourced sensor data with street-level imagery to uncover previously invisible risk patterns [18]. The present study builds on that trajectory by focusing on the full-system perspective of deploying such models as urban informatical infrastructure.

3. System Architecture and Data Integration

Designing a deep learning-based bicycle crash risk modeling system demands an architectural perspective that spans data acquisition, preprocessing, model inference, and feedback loops. The system must ingest a heterogeneous collection of data streams, each with distinct sampling rates, spatial coverages, error profiles, and privacy constraints. Mobile sensing data from cyclists can include continuous GPS trajectories, accelerometer readings, gyroscope orientations, and optionally physiological indicators from wearable devices. These high-frequency kinematic feeds capture the dynamic behavior of individual cyclists, but they are inherently voluntary and unevenly distributed across the population. The spatial coverage bias introduced by app-based volunteering raises immediate questions about representativeness and the transferability of risk models across demographic and geographic segments.

Complementing the mobile streams are urban built environment datasets typically maintained by municipal planning departments and open data portals. These encompass road geometry, intersection types, bicycle lane presence, lane widths, traffic signal configurations, speed limits, land use classifications, tree canopy coverage, streetlight locations, and transit stop densities. The temporal refresh rates of these layers vary dramatically: while street network updates occur on multi-year planning cycles, land use and zoning changes may be annual, and traffic signal timing plans can be adjusted seasonally or dynamically. The modeling architecture must therefore accommodate the fusion of fast-moving observational data with slow-changing contextual features, a requirement that has profound implications for model training, retraining, and drift detection. In vehicular crash prediction, the spatial units of analysis are often fixed segments or intersections, but for bicycles, the risk surface is continuous and path-dependent, influenced by the sequence of environments a cyclist traverses. Representing the spatial domain as a dynamic graph, in which nodes correspond to intersections or mid-block segments and edges capture connectivity and risk propagation, offers a flexible abstraction that can be refined over time as data density increases.

The data pipeline must also address the fundamental tension between personal privacy and model fidelity. Raw GPS traces can uniquely identify individuals even after simple pseudonymization [12]. Aggregation techniques such as differential privacy, spatial k-anonymity, and on-device federated learning emerge as necessary architectural components rather than afterthoughts. Federated learning, where model weights are updated locally on users' devices and only aggregated gradients are shared with a central server, preserves raw data locality and reduces privacy risks, but it introduces additional complexity in terms of non-independent and identically distributed data partitions and communication overhead. The system architecture must therefore make explicit trade-offs among model accuracy, update latency, energy consumption on mobile devices, and privacy guarantees. Additionally, the governance of data access across multiple stakeholders—including transportation agencies, public health departments, private mobility operators, and academic researchers—requires interoperable application programming interfaces, standardized metadata schemas, and clearly defined data stewardship roles that align with municipal data governance policies.

4. Deep Learning Model Design and Spatiotemporal Structure

The core predictive engine of the system rests on a spatiotemporal deep learning model that accepts multivariate inputs defined over a spatial graph and a continuous time axis. Unlike count-based crash prediction models that discretize space into regular grid cells or fixed-length road segments, bicycle risk modeling benefits from irregular graph representations where spatial nodes represent functional units such as intersections, crosswalks, and bike lane segments. Graph convolutional networks propagate information across these nodes according

to learned adjacency weightings that can incorporate not only topological connectivity but also similarity metrics derived from built environment features [6]. The temporal dimension is equally critical, as risk fluctuates at multiple scales: diurnal commuting peaks, weekly recreational patterns, seasonal weather effects, and long-term infrastructure changes. Attention-driven sequence models, particularly transformer variants, provide the capability to attend over variable-length historical windows and learn which past observations are most diagnostic of current risk conditions [7]. In the context of bicycle safety, an attention mechanism might learn to heavily weight a sudden surge in average cyclist deceleration rates observed ten minutes prior on an upstream segment when predicting near-miss likelihood at a downstream intersection.

The integration of built environment features as static or slowly varying node attributes raises the question of how to best encode them within deep architectures. One approach treats land use, bike facility type, and intersection control as categorical embeddings learned jointly with the spatiotemporal model, allowing the network to discover interactions between these context features and dynamic kinematic patterns. Another pathway, inspired by multimodal learning, employs separate encoder branches for static imagery such as street-level photos and satellite views, then fuses the resulting visual representations with kinematic and structural graph data. The choice of fusion strategy—early fusion at the input level, late fusion at the decision level, or intermediate attention-gated fusion—has downstream effects on model interpretability and modularity. Early fusion can hinder the ability to diagnose which modality drives a particular risk prediction, whereas intermediate fusion with attention mechanisms can produce modality-specific saliency maps that aid in explaining model outputs to transportation planners. This interpretability is not merely an academic nicety; it is a prerequisite for generating actionable recommendations that can be linked to specific infrastructure modifications, such as prioritizing signal retiming or installing protected bike lanes at high-risk locations identified by the model.

Another architectural consideration hinges on the handling of rare events. Bicycle crashes are statistically rare relative to the volume of non-crash observations, creating extreme class imbalance that can lead deep learning models to collapse into trivial non-crash predictions. Specialized loss functions, such as focal loss or ranking-based objectives, alongside data augmentation strategies that synthesize near-miss scenarios through kinematic perturbation models, must be structurally embedded in the training pipeline. Furthermore, the evaluation of model performance cannot be reduced to global metrics such as area under the receiver operating characteristic curve alone; rather, spatially and temporally stratified evaluation that measures predictive equity across neighborhoods of different income levels, racial compositions, and cycling infrastructure densities is essential to prevent models from systematically underestimating risk in underserved areas. The calibration of predicted probabilities is another system-level concern, since risk scores used in municipal resource allocation must faithfully reflect absolute risk magnitudes to avoid misallocation of limited safety improvement funds.

5. Deployment, Governance, and Infrastructure Considerations

Transitioning a deep learning bicycle crash risk model from a research prototype to a sustained operational service involves embedding the model within an urban informatics infrastructure that interfaces with existing transportation management systems. This infrastructure must support continuous data ingestion, automated model retraining to combat concept drift as cycling volumes and infrastructure change, and the dissemination of risk

heatmaps through web dashboards, application programming interfaces, and push notifications to navigation applications. The frequency of model updates introduces a tension between freshness and computational cost: inferring risk at sub-hourly resolution for an entire city requires scalable inference pipelines that can leverage cloud or edge computing resources. Edge computing, where inference is partially executed on roadside units or even on smart bicycles themselves, offers latency advantages and reduces data transmission volumes, but it constrains model complexity due to hardware limitations and complicates version consistency across a distributed network.

Governance mechanisms must delineate responsibilities for data curation, model monitoring, and accountability for decisions influenced by risk predictions. In many cities, transportation safety is a distributed responsibility spanning departments of transportation, public works, police, and public health. A centralized predictive model can easily become an orphaned asset without clear institutional ownership and a sustainability plan for its maintenance. Drawing from socio-technical transition theory, the introduction of an artificial intelligence-enabled safety system can be viewed as a niche innovation that must navigate the existing regime of established engineering standards, liability norms, and professional practices [15]. For instance, traffic engineers accustomed to warrants based on historical crash counts may be skeptical of data-driven risk scores that cannot be easily linked to a specific empirical crash record. Building institutional trust requires not only transparent model architectures but also longitudinal validation studies that track the correspondence between predicted risk hotspots and subsequent safety outcomes, ideally within a controlled quasi-experimental framework such as before-after comparisons with matched control sites.

The interoperability of the system with existing data standards, such as the General Transit Feed Specification for transit data or the Mobility Data Specification for shared micromobility, can accelerate adoption by reducing integration friction. However, this standardization must be balanced against the need for local customization: the relative importance of built environment variables, the cultural norms of cycling behavior, and the typology of bicycle infrastructure vary markedly across global cities. A model trained predominantly on data from Copenhagen or Amsterdam may not transfer well to Jakarta or Los Angeles without substantial domain adaptation. Thus, the deployment architecture should support transfer learning and fine-tuning workflows that allow a base model to be adapted to a new city with a fraction of the labeled crash data. This modular approach also enables incremental incorporation of locally specific data sources, such as community-reported hazard inventories or municipal maintenance records, without requiring a wholesale model replacement.

6. Fairness, Equity, and Ethical Implications

Deep learning models used to inform public safety resource allocation inevitably raise questions of distributive justice. Bicycle crash risk is not distributed uniformly across urban populations; it is shaped by historical patterns of infrastructure investment that often favored automobile-oriented suburbs over dense, lower-income urban cores where cycling mode shares may be higher yet facilities are scarcer. If a risk model is trained on observed crash and near-miss data that reflect current cycling volumes, it may inadvertently direct resources toward areas with high absolute crash counts, which tend to be affluent areas with high cycling participation, while neglecting neighborhoods where suppressed demand and safety fears deter cycling altogether. This feedback loop, analogous to the "selective labeling" problem in machine learning fairness literature, can perpetuate existing infrastructure inequities [14]. The system design must actively counteract such dynamics by incorporating

equity-weighted objectives, such as prioritizing risk reductions in historically underserved census tracts or applying subgroup-specific calibration that ensures predictive accuracy is not systematically lower for minority cyclists.

The ethical framework within which such systems operate must go beyond technical fairness metrics to embrace a broader principle of procedural justice. Communities affected by the deployment of automated risk assessment tools should have meaningful opportunities to participate in defining the problem formulation, selecting outcome variables, and setting the thresholds that trigger resource allocation. The governance architecture should include multi-stakeholder advisory boards that represent cyclists, advocacy groups, city planners, and privacy advocates, with explicit mandates to audit model performance and contest decisions. Algorithmic impact assessments, similar to environmental impact assessments, are emerging as a governance instrument that could be applied to predictive safety systems, requiring developers and deploying agencies to document potential disparate impacts, data provenance, and fallback mechanisms before full-scale deployment [16]. Such processes can surface value conflicts early, such as the tension between maximizing total crash reductions versus equalizing risk reductions across neighborhoods, and force transparent deliberation rather than embedding contested values in opaque model parameters.

The interpretability of model outputs also intersects with fairness when risk scores are used to justify infrastructure investments that have long-term lock-in effects. A black-box model that identifies an intersection as high-risk without revealing whether the signal is driven by inadequate sight distance, insufficient crossing time, or aggressive motorist behavior provides little guidance for designing an effective countermeasure. Attention maps and feature attribution techniques can partially illuminate these pathways, but they must be communicated through careful design of decision-support interfaces that do not overwhelm non-technical stakeholders with false precision. Fairness-aware model design thus encompasses not only the mathematical definition of fairness constraints but also the human-computer interaction design of explanation interfaces that support equitable decision-making processes.

7. Robustness, Sustainability, and Resilience

Any safety-critical system that informs real-world infrastructure decisions must be robust to distributional shifts, data quality degradation, and adversarial manipulation. Urban environments are non-stationary; cycling patterns changed dramatically during the COVID-19 pandemic, with temporary bicycle lanes, changed commuting patterns, and a surge in recreational cycling altering the underlying relationship between built environment and risk [1]. A model that does not detect and adapt to such shifts can produce miscalibrated risk estimates that erode user trust and potentially misallocate safety resources. Robustness engineering thus requires continuous monitoring of input feature distributions, output score drift, and ground-truth crash outcomes using statistical process control techniques. When drift is detected, automated retraining pipelines—or in more conservative configurations, human-in-the-loop review processes—must be triggered. The sustainability of such monitoring infrastructures depends on securing long-term funding and staffing beyond initial research grants, a challenge that has plagued many smart city initiatives. Business models that embed the cost of model maintenance within municipal capital improvement programs or public health budgets can provide more durable support than innovation-grant cycles.

The resilience of the system to data interruptions is another crucial facet. Mobile sensing data streams can be disrupted by application platform policy changes, privacy regulations, or shifts

in user participation. The architecture must therefore be designed to gracefully degrade, for instance by falling back to coarser spatial models when fine-grained kinematic data are sparse, or by leveraging transfer from motor vehicle sensor networks to maintain risk coverage in data-poor areas. The concept of infrastructure resilience, originally applied to physical networks such as transportation and telecommunications, is increasingly relevant to data-driven decision systems [22]. Building redundancy into the sensing fabric—combining dedicated cycling sensors with opportunistic data from navigation apps, dockless bikeshare systems, and third-party aggregators—enhances resilience and reduces single points of failure. However, such multi-source integration must navigate commercial data sharing agreements, licensing terms, and competitive dynamics among mobility providers, which often resist open data sharing for business advantage.

The long-term sustainability of deep learning-based safety systems also intersects with environmental sustainability. Training large-scale transformer models entails substantial energy consumption and carbon emissions. Edge inference and model compression techniques such as quantization and pruning can reduce the operational carbon footprint, but the fundamental question remains whether the net social benefit of highly granular risk prediction justifies the computational resources when simpler statistical models might suffice for many use cases. A responsible deployment strategy should include a tiered model portfolio: lightweight models for continuous city-wide risk monitoring, and heavier, more expressive models for in-depth post-hoc analysis of specific corridors or before-after evaluations. This stratified approach mirrors the tiered architecture of modern web services and aligns resource consumption with the criticality and latency of the decision context.

8. Policy and Planning Implications

The integration of deep learning bicycle crash risk models into municipal planning processes has the potential to fundamentally reshape the safety management lifecycle. Instead of relying solely on reactive crash hot spot identification that typically requires three to five years of post-implementation data to evaluate, planners could use predictive risk surfaces to prioritize candidate projects, simulate the safety impacts of alternative designs before construction, and monitor the evolution of risk as land use and travel behavior change. This shift toward anticipatory governance would represent a significant departure from prevailing norms, requiring updates to engineering guidelines that currently specify minimum crash history thresholds for trigger-based interventions like traffic signal installation or road diets. Cities that have pioneered Vision Zero commitments are natural early adopters, as their policy objectives demand proactive and systematic elimination of traffic fatalities rather than incremental hot spot fixes [23].

The policy levers enabled by predictive risk models extend beyond capital infrastructure projects. Dynamic interventions, such as temporarily lowering speed limits, activating in-pavement lighting at crosswalks, or rerouting bike share fleets away from high-risk areas during nighttime or adverse weather, can be triggered algorithmically based on real-time risk forecasts. These interventions, however, sit in a regulatory gray zone, as existing traffic control device manuals have not kept pace with the possibility of algorithmically dynamic traffic control. The institutional infrastructure for certification, testing, and liability attribution for such adaptive systems remains underdeveloped. Policy frameworks must evolve in tandem with technical capabilities, creating sandbox environments where municipalities can experiment with adaptive safety measures under close monitoring, with clear mechanisms for public notice and redress when automated interventions prove intrusive or counterproductive.

The data governance policies required for bicycle crash risk modeling also intersect with broader societal debates about surveillance and the public sphere. While aggregated mobility data provide immense public value, the continuous collection of location traces, even with consent, can contribute to a chilling effect on cycling participation if the public perceives that their movements are being monitored by government agencies. Transparent data minimization protocols, purpose limitation clauses, and independent oversight bodies are necessary to maintain public trust. Lessons from the smart city backlash in Toronto's Quayside project highlight the perils of progressing with technologically ambitious urban sensing without meaningful democratic deliberation [24]. Bicycle safety systems must therefore be codesigned with privacy commissioners and civil liberties organizations, not merely presented as a *fait accompli* after technical development. Establishing clear sunset clauses for data retention and strong prohibitions against repurposing safety data for law enforcement or immigration enforcement is essential to safeguard the civic nature of the infrastructure.

9. Conclusion

This paper has examined the design, deployment, and governance of deep learning-based spatiotemporal modeling systems for bicycle crash risk that leverage the convergence of mobile sensing and rich urban built environment data. The shift from reactive, aggregate crash analysis to proactive, high-resolution risk surface estimation holds transformative potential for urban safety planning, but realizing this potential depends on rigorous attention to a constellation of system-level considerations that extend far beyond model architecture. The architectural choices of graph representation, attention modulation, multimodal fusion, and rare-event handling structure the baseline predictive capability, yet the effectiveness and legitimacy of such systems are equally determined by the data fusion pipelines that preserve privacy, the deployment infrastructures that ensure robustness and equity, and the governance frameworks that embed algorithmic outputs within democratic planning processes. Cross-domain insights from traffic flow prediction, crash analysis on motorways, and urban computing illuminate the technical pathways, while literature on socio-technical transitions, algorithmic fairness, and urban resilience underscores the institutional and ethical imperatives.

As cities accelerate their adoption of digital twins and intelligent transportation systems, bicycle safety must not be relegated to an afterthought retrofitted onto car-centric sensor networks. Instead, the unique kinematic fragility, environmental sensitivity, and equity dimensions of cycling demand bespoke modeling ecosystems that treat cyclists as first-class sensing agents rather than passive risks. Future research must advance the empirical validation of these systems through multi-city longitudinal studies that track not only predictive accuracy but also downstream safety outcomes, infrastructure investment patterns, and changes in cycling participation across diverse demographic groups. Only through such holistic evaluation can the academic community and its municipal partners distinguish between models that merely publish and those that genuinely protect. The infrastructure of deep learning for bicycle safety is not only a technical artifact but a socio-technical commitment to a more equitable, sustainable, and humane urban future.

References

1. Pucher, J., Dill, J., & Handy, S. (2010). Infrastructure, programs, and policies to increase bicycling: an international review. *Preventive Medicine*, 50, S106-S125.

2. Beck, L. F., Dellinger, A. M., & O'Neil, M. E. (2007). Motor vehicle crash injury rates by mode of travel, United States: using exposure-based methods to quantify differences. *American Journal of Epidemiology*, 166(2), 212-218.
3. Ewing, R., & Cervero, R. (2010). Travel and the built environment: A meta-analysis. *Journal of the American Planning Association*, 76(3), 265-294.
4. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
5. Zhang, J., Zheng, Y., & Qi, D. (2017). Deep spatio-temporal residual networks for citywide crowd flows prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1).
6. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*.
7. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
8. Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., & Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. *Journal of Big Data*, 2(1), 1-21.
9. Bao, J., Liu, P., Yu, H., & Xu, C. (2019). Incorporating temporal correlation into a deep learning framework for short-term crash occurrence prediction on expressways. *Accident Analysis & Prevention*, 131, 248-256.
10. Mannering, F. L., & Bhat, C. R. (2014). Analytic methods in accident research: Methodological frontier and future directions. *Analytic Methods in Accident Research*, 1, 1-22.
11. Dozza, M., & Werneke, J. (2014). Introducing naturalistic cycling data: What factors influence bicyclists' safety? *Transportation Research Part F: Traffic Psychology and Behaviour*, 24, 83-91.
12. de Montjoye, Y. A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013). Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, 1376.
13. 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.
14. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104, 671-732.
15. Geels, F. W. (2002). Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. *Research Policy*, 31(8-9), 1257-1274.
16. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707.
17. Wang, Z., He, S. Y., & Leung, Y. (2018). Applying mobile phone data to travel behaviour research: A literature review. *Travel Behaviour and Society*, 11, 141-155.

18. Li, Y., Zhang, Y., Long, Y., Bhalla, K., & Ezzati, M. (2025). Assessing bicycle safety risks using emerging mobile sensing data. *Travel Behaviour and Society*, 38, 100906.
19. Zahabi, S. A. H., Strauss, J., Manaugh, K., & Miranda-Moreno, L. F. (2011). Estimating potential effect of speed limits, built environment, and other factors on severity of pedestrian and cyclist injuries in crashes. *Transportation Research Record*, 2247(1), 81-90.
20. Yao, H., Tang, X., Wei, H., Zheng, G., & Li, Z. (2019). Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 5666-5673.
21. Zou, Y., Wang, X., & Zhang, D. (2019). Examining the factors influencing pedestrian and cyclist crash frequency by severity using the Conway–Maxwell–Poisson model. *Accident Analysis & Prevention*, 132, 105284.
22. Wang, J., & Muddassir, A. B. M. (2014). A review of resilience in transportation systems. *Transport Reviews*, 34(2), 183-208.
23. Pucher, J., & Buehler, R. (2008). Making cycling irresistible: lessons from the Netherlands, Denmark and Germany. *Transport Reviews*, 28(4), 495-528.
24. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.
25. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.