Research Statement

Bhumika

Ph.D. Candidate, Department of Computer Science & Engineering, IIT Jodhpur
Email: bhumika.1@iitj.ac.in | Google Scholar

Urbanization is rapidly accelerating worldwide, driven by scientific and technological advancements. Ensuring livability and sustainable development in cities requires the ability to forecast and manage **urban incidents** such as crime, accidents, stempede, flood, epidemic, etc. proactively. My Ph.D. research, titled *"Spatio-Temporal Incident Analytics in Cityspaces,"* focuses on identifying and forecasting these impactful events in urban environments using real-time and historical data. With the proliferation of sensing and logging technologies, large-scale urban data—ranging from traffic patterns to environmental readings—has become increasingly accessible. By leveraging deep learning techniques in conjunction with edge computing, my work aims to extract actionable insights for applications such as crime prediction, accident detection, and environmental quality monitoring. The ultimate goal is to support city administrators and policymakers in making data-driven decisions that enhance public safety, operational efficiency, and quality of life. My research lies at the intersection of **spatial and temporal learning** with a strong focus on **AI for social good**.

Research Thrust 1: Spatio-Temporal Learning for Citywide Incident Prediction

In my Ph.D. research, I focus on developing robust learning frameworks for citywide incident prediction by capturing spatio-temporal-contextual patterns from urban data that is inherently sparse, heterogeneous, and dynamic. One of the core challenges in this domain is data sparsity—particularly for rare or region-specific events. To address this, I incorporated rich contextual signals into the learning process, such as: Point-of-interest (POI) distributions to model functional dependencies between regions, Weather conditions (e.g., temperature, rainfall, cloud cover) to reflect environment-induced behavioral shifts, Temporal cues including calendar events, weekends, and diurnal cycles, and Taxi flow patterns as a proxy for population mobility and density. This approach led to the development of a spatio-temporal deep learning framework for citywide incident forecasting, which was published at ECML-PKDD 2022 and selected for presentation in the Ph.D. Forum of the same conference.

Building on this, I also explored risk-aware route recommendations through predictive modeling. In this work, I designed a multitask transformer architecture capable of forecasting future traffic congestion, accident likelihood, and crime risk to recommend safer and more efficient travel routes. This work was accepted at ACM RecSys 2022, highlighting the practical applicability of spatio-temporal learning in real-time urban mobility systems.

Research Thrust 2: Federated Perspective of Citywide Incident Prediction.

To address the challenges of data decentralization, regional heterogeneity, and privacy preservation in citywide incident prediction, I extended my research into the domain of federated and edge-based learning architectures. In our proposed framework, we introduce a hybrid learning architecture comprising edge-level local computation, centralized server aggregation, and a novel peer-to-peer collaborative learning mechanism. The peer-to-peer layer leverages the principle of spatial locality, allowing geographically proximate or contextually similar regions to influence each other's learning process, which is critical for capturing neighborhood-level patterns in urban environments.

One major challenge in this setting is quantity and label heterogeneity—where different regions exhibit vastly different incident frequencies and class distributions. A single global model often fails to generalize well across such non-iid data. To overcome this, we transitioned to a federated learning paradigm, enabling each regional client to retain autonomy in model updates while still contributing to a shared global model. This work was carried out in collaboration with Prof. Philippe Lalanda during my research visit to the University of Grenoble Alpes, France.

To further enhance the robustness of aggregation, we proposed a personalized federated learning approach that integrates graph convolutional operations at the server level, allowing parameters from structurally and behaviorally similar regions to be more closely aligned during model aggregation. This reduces the negative influence of highly dissimilar clients and improves learning performance in skewed urban settings.Additionally, we identified and addressed the zero-inflation problem, where certain regions experience very few incident occurrences, leading to biased model behavior. To mitigate this, we modified the learning objective using negative binomial distribution-based loss functions, which better accommodate sparse event data and improve prediction accuracy in low-activity regions.

Research Thrust 3: Bias Handling in Citywide Incident Prediction.

Bias in machine learning models often emerges from imbalanced data distributions, where certain classes or regions are underrepresented. In the context of citywide incident prediction, this bias is typically manifested as a spatial long-tail distribution—a scenario in which a small number of regions (e.g., city centers) generate abundant incident data, while the majority of regions (e.g., suburban or low-density areas) contribute very few data points. This imbalance leads to prediction bias, where models perform well on data-rich areas but poorly on data-scarce regions, thereby reinforcing spatial inequality in model outcomes.

To address this critical challenge, I focused on developing spatial long-tail learning techniques aimed at improving model generalization across both data-rich and data-sparse regions. By improving the performance of predictive models on underrepresented areas, this research advances the goal of inclusive and responsible AI for smart cities.

Future Directions and Outlook

Building upon my doctoral work, I plan to extend my research along several key directions that lie at the intersection of trustworthy, responsible, privacy preserving and adaptable AI. These directions address the growing need for robust, adaptive, and socially responsible machine learning models in real-world urban systems:

1. **Adversarial Robustness in Spatio-Temporal GNNs**: I aim to investigate the vulnerability of dynamic graph neural networks to adversarial attacks, particularly in the presence of spatially long-tailed and imbalanced data. This is a relatively unexplored area, and I intend to design robust architectures that can detect and withstand such attacks while maintaining predictive accuracy for both high- and low-incident regions.
2. **Responsible and Explainable AI in Urban Systems**: To foster transparency and trust in AI-driven decision-making, I intend to develop interpretable, fairness-aware models that provide clear explanations for predictions related to mobility, crime, and accident. This aligns with the growing emphasis on accountability and contestability in AI used for public infrastructure and governance.
3. **Continual Learning in Urban Systems**: I plan to integrate continual learning mechanisms into urban AI frameworks, enabling models to dynamically update in response to periodic influxes of new data. This includes designing triggers for model retraining, mechanisms for seamless model integration, and strategies to ensure minimal disruption—all with the goal of maintaining accuracy, reducing manual intervention, and enhancing responsiveness to evolving urban conditions.
4. **Cross-city Knowledge Transfer**: I aim to enable transfer learning across cities or tasks where labeled data is scarce, particularly for rare but critical incidents such as disasters or criminal activities. This includes developing generalizable representations that adapt across domains while preserving region-specific nuances, thus improving model utility in data-sparse or under-resourced urban areas.

My research has led to the release of scripts for data extraction and different models for prediction. Different works have resulted in over ten peer-reviewed publications in CORE A/B-ranked venues (RecSys, ECML-PKDD, ITSC, VTC, IJCNN), along with a published Indian patent.