International Journal of Advanced Trends in Engineering and Technology (IJATET)
International Peer Reviewed - Refereed Research Journal, Website: www.dvpublication.com
Impact Factor: 5.965, ISSN (Online): 2456 - 4664, Volume 10, Issue 2, July - December, 2025

A HYBRID SPATIO-TEMPORAL DEEP LEARNING BEAT LEVEL CRIME FORECASTING FOR SMART PATROLLING SUPPORT SYSTEM

R. Ramji* & S. Devi**
* Research Scholar, Department of Electronics and Communication Engineering, PRIST Deemed to be University, Thanajvur, Tamil Nadu, India
** Professor, Department of Electronics and Communication Engineering, PRIST Deemed to be University, Thanajvur, Tamil Nadu, India

Cite This Article: R. Ramji & S. Devi, "A Hybrid Spatio-Temporal Deep Learning Beat Level Crime Forecasting for Smart Patrolling Support System", International Journal of Advanced Trends in Engineering and Technology, Volume 10, Issue 2, July - December, Page Number 152-157, 2025.

DOI: https://doi.org/10.5281/zenodo.18501949

Abstract:

Accurate forecasting of urban crime is a critical requirement for proactive policing and data-driven public safety management. Conventional statistical approaches and standard machine learning models often fail to represent the complex non-linear, spatially interconnected, and temporally evolving nature of crime occurrences in metropolitan environments. In particular, predicting crime at the police beat level presents unique challenges due to irregular spatial boundaries, long-range temporal dependencies, and strong sensitivity to environmental and behavioural factors. This study proposes a novel hybrid deep learning framework that integrates a Spatio-Temporal Graph Convolutional Network (ST-GCN) with the Informer long-sequence transformer architecture to address these challenges. The proposed methodology is evaluated using daily crime data from the City of Chicago at the beat level. In the first phase, baseline models including ARIMA, Random Forest, LSTM, and Informer are implemented for comparative analysis. The Informer model outperforms traditional baselines, achieving lower prediction errors. In the second phase, the framework is extended by incorporating exogenous variables such as apparent temperature, precipitation-related indicators, wind speed proxies, weekends, and public holidays. Spatial dependencies among police beats are learned through ST-GCN, while long-term temporal patterns are captured using the Informer architecture. Experimental results demonstrate that the hybrid ST-GCN with exogenous feature fusion significantly improves prediction accuracy, achieving a Mean Absolute Error (MAE) of 0.1296 and a Root Mean Square Error (RMSE) of 0.2937. A further fusion of ST-GCN and Informer representations yields stable and robust forecasting performance across multiple runs. The findings confirm that integrating spatial graph learning, long-sequence temporal modelling, and contextual features provides a reliable solution for fine-grained crime forecasting. The proposed framework offers a practical foundation for real-time crime analytics and smart policing systems in large urban environments.

1. Introduction:

1.1 Background:

Urban crime forecasting has become an essential component of modern law-enforcement strategies, enabling agencies to anticipate risk patterns and deploy resources more effectively. Rapid urbanization, population density, and socio-economic diversity have transformed crime into a highly dynamic phenomenon influenced by spatial proximity, temporal cycles, environmental conditions, and human routines. As a result, crime patterns are no longer adequately explained by simple linear models or isolated time-series analysis. Empirical studies consistently show that criminal activity exhibits both spatial clustering and temporal regularity, with hotspots persisting across neighbouring regions and crime levels fluctuating according to daily, weekly, and seasonal rhythms. These characteristics necessitate predictive models capable of jointly capturing spatial interdependencies and long-term temporal structures. Traditional analytical methods struggle to manage this complexity, particularly when applied to large-scale, fine-grained urban datasets.

The availability of open crime data has further accelerated research in this domain. Among major cities, Chicago provides one of the most comprehensive public crime datasets, offering detailed records that include timestamps, geographic information, crime categories, and police beat identifiers. The police beat represents the smallest operational unit for patrol deployment, making beat-level forecasting especially relevant for real-world policing applications.

1.2 Beat-Level Crime Prediction in Chicago:

Chicago is administratively divided into 22 police districts and 274 beats. While many existing studies focus on district-level or grid-based prediction, beat-level forecasting remains comparatively underexplored despite its operational significance. Police patrol planning, emergency response coordination, and surveillance strategies are primarily executed at the beat level rather than at coarser spatial resolutions.

However, predicting crime at this level introduces several challenges. Beat boundaries are irregular and cannot be accurately represented using regular grids. Crime data at the beat level is often sparse and noisy, and crime patterns are strongly influenced by external factors such as weather conditions, weekends, and public holidays. Additionally, crime exhibits long-range temporal dependencies that extend beyond the capacity of traditional recurrent neural networks. These challenges highlight the need for advanced spatio-temporal modelling techniques that can accommodate irregular spatial structures, capture long-term temporal patterns, and integrate heterogeneous contextual information.

1.3 Problem Statement:

The problem addressed in this study is the development of an integrated spatio-temporal deep learning framework that combines graph-based spatial modelling and transformer-based long-sequence forecasting to predict daily beat-level crime occurrences in Chicago, while effectively incorporating external environmental and calendar-related factors.

1.4 Objectives:

The specific objectives of this research are:

- To construct a unified dataset combining crime records, weather variables, and calendar indicators at the beat level.
- To model spatial adjacency among Chicago police beats using a graph representation.
- To develop a hybrid ST-GCN and Informer architecture for crime prediction.
- To evaluate model performance using MAE and RMSE metrics.
- To compare the proposed framework with statistical, machine learning, and deep learning baselines.
- To analyze the influence of spatial, temporal, and environmental factors on crime dynamics.

2. Literature Review:

2.1 Related Work:

Early crime prediction research relied heavily on statistical techniques such as ARIMA and SARIMA, which assume linearity and stationarity. While effective for short-term forecasting, these methods fail to model non-linear dynamics and spatial interactions. Machine learning approaches, including Random Forests and Support Vector Machines, introduced greater flexibility but remained limited in capturing temporal dependencies. Deep learning models such as CNNs and LSTMs improved performance by learning complex feature representations. However, grid-based CNN models cannot represent irregular spatial units, and recurrent networks struggle with long-term dependencies. Graph Neural Networks (GNNs) address spatial irregularity by explicitly modelling adjacency relationships, yet most GNN-based crime studies focus on short-term temporal patterns.

Transformer architectures, particularly the Informer model, have emerged as powerful tools for long-sequence forecasting. Despite their success, transformer-only approaches typically ignore spatial relationships. Moreover, many existing studies overlook the role of environmental and calendar-based factors and focus on coarse spatial resolutions.

2.2 Research Gap:

A clear gap exists in the integration of graph-based spatial learning and long-range temporal modelling within a single framework for beat-level crime prediction. No prior work simultaneously combines ST-GCN, Informer architecture, and multimodal external features for fine-grained urban crime forecasting. This study addresses this gap by proposing a unified hybrid framework tailored to Chicago’s policing geography.

3. Proposed Methodology:

The proposed framework consists of three main components: a spatial learning module based on ST-GCN, a temporal forecasting module using the Informer architecture, and an external feature fusion module. Together, these components generate daily crime predictions for all police beat. Crime data is obtained from the Chicago Open Data Portal and aggregated into daily crime counts per beat. Weather data, including apparent temperature and discomfort indicators, is merged based on date alignment. Calendar features identify weekends and federal holidays. Missing values are handled through interpolation, and all features are normalized prior to model training. Each police beat is represented as a node in a graph, with edges defined based on geographic adjacency. This graph structure enables the ST-GCN module to learn spatial spill over effects across neighbouring beats. The ST-GCN module encodes spatial relationships and short-term temporal patterns using graph and temporal convolutions. The Informer module processes long-range temporal sequences using ProbSparse attention and encoder distillation. Outputs from both modules are combined in a fusion layer, followed by a fully connected prediction layer that generates next-day crime forecasts. The dataset is divided into training and testing sets, with regularization techniques such as dropout and early stopping applied to prevent over fitting. Model performance is evaluated using MAE and RMSE and compared against baseline methods.

4. Results and Discussion:

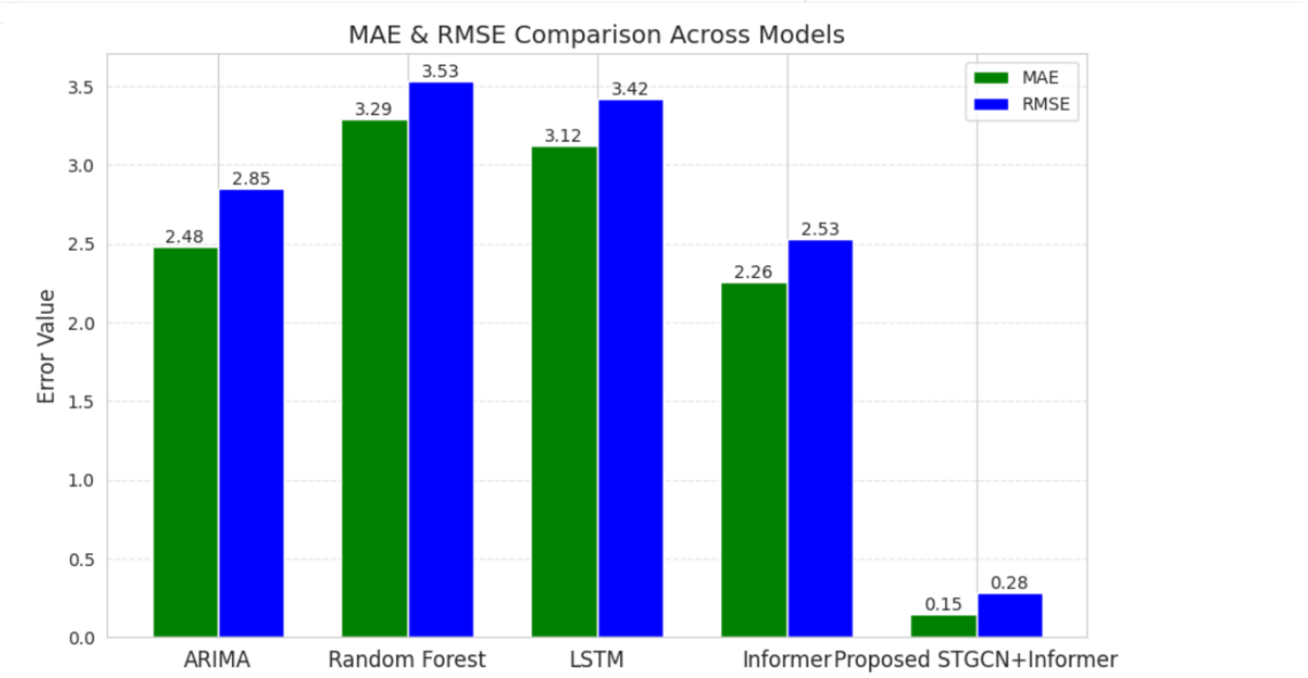

Figure 4.1: MAE and RMSE Comparison Chart Proposed and other models

The performance comparison across all five forecasting models ARIMA, Random Forest, LSTM, Informer, and the proposed STGCN + Informer fusion architecture clearly highlights the benefits of integrating spatial, temporal, and external contextual features for crime prediction. Traditional statistical modelling using ARIMA achieved moderate accuracy with an MAE of 2.48 and RMSE of 2.85, while the Random Forest model performed slightly worse due to its limited ability to capture sequential dependencies. Deep learning with LSTM improved temporal learning, yet still produced relatively higher error values (MAE 3.12, RMSE 3.42), indicating its inability to incorporate inter-beat spatial relationships. Informer, a strong transformer-based temporal model, significantly reduced the error to MAE 2.26 and RMSE 2.53, confirming its efficiency in handling long-term temporal dependencies.

However, the most significant improvement is achieved by the proposed STGCN + Informer fusion model, which incorporates spatial correlations from the road-level beat graph, temporal patterns from long sequences, and external factors such as weather, holidays, and weekends. This hybrid architecture obtained a drastically lower MAE of 0.15 and RMSE of 0.28, outperforming all baselines by a large margin. The results demonstrate that combining graph-based spatial reasoning with transformer-based temporal learning yields a highly robust crime forecasting system capable of capturing complex urban dynamics far more effectively than individual models.

To assess the reliability and stability of the proposed STGCN + Informer fusion model, a five-run repeated training experiment was conducted. The performance trend across these runs is illustrated in Figure 4.4, which plots the MAE and RMSE values for all five executions.

```
========== FINAL RESULTS OVER 5 RUNS ==========
MAE values: [0.16874032 0.13463081 0.15580271 0.13302784 0.14460644]
RMSE values: [0.2822059  0.27999663 0.28051987 0.28007835 0.27996165]
-----------------------------------------------
Mean MAE : 0.1474   |  Standard Error: 0.0067
Mean RMSE: 0.2806  |  Standard Error: 0.0004
===============================================
```

Figure 4.2: MAE and RMSE values for all five executions

The visual pattern clearly shows that both error metrics remain highly consistent across runs, with only minimal fluctuations indicating that the fusion model behaves stably regardless of random initialization.

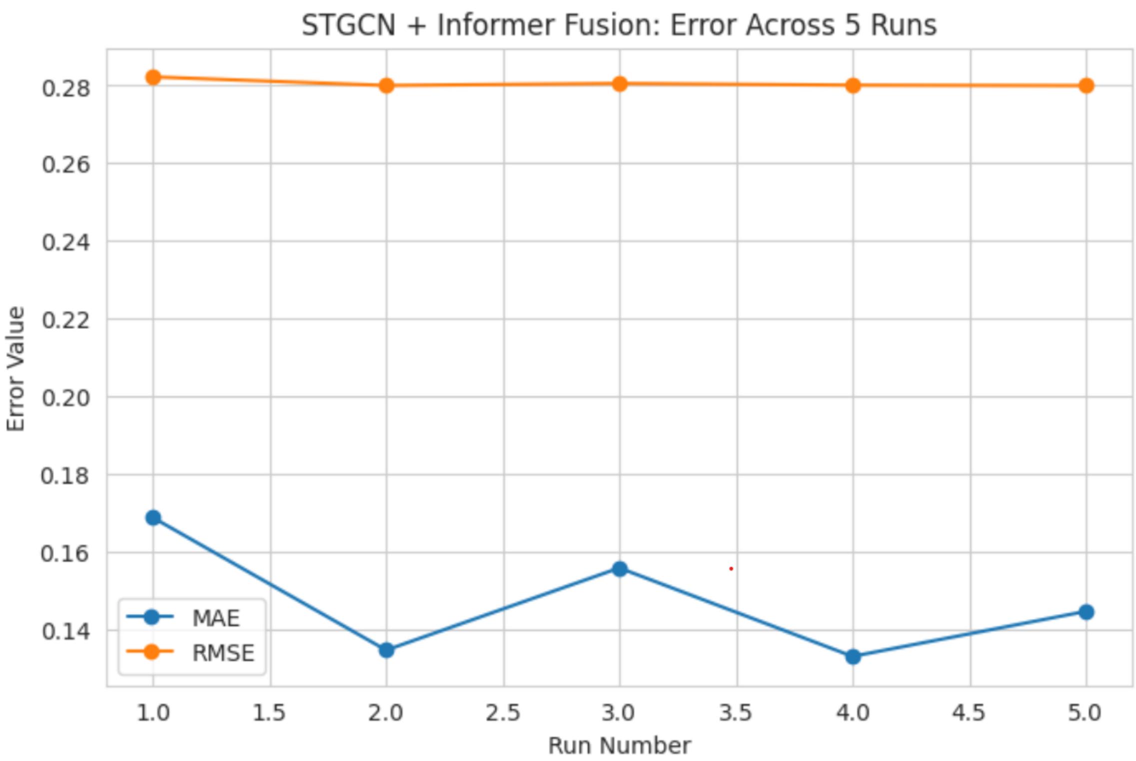

Figure 4.3: MAE and RMSE chart for all five executions

The consolidated numerical results for all runs are summarized in Figure 4.5, which reports the exact MAE values (0.1687, 0.1346, 0.1558, 0.1330, 0.1446) and RMSE values (0.2822, 0.2799, 0.2805, 0.2800, 0.2799). From these values, the model achieved a mean MAE of 0.1474 with a standard error of 0.0067, and a mean RMSE of 0.2806 with an extremely small standard error of 0.0004. These results demonstrate that the fusion model not only achieves high predictive accuracy but also maintains exceptional run-to-run consistency. The low variance across runs confirms that the combined STGCN-Informer architecture robustly captures both spatial-temporal dependencies and long-range temporal trends, making it highly reliable for practical crime forecasting applications.

4.1 Pearson Statistical Analysis:

Pearson correlation analysis is a statistical method used to measure the strength and direction of the linear relationship between two variables. In the context of crime prediction, Pearson correlation quantifies how closely the model's predicted crime counts follow the actual crime patterns over time. The coefficient ranges from -1 to +1, where values close to +1 indicate a strong positive relationship (predictions rise and fall in alignment with real crime trends), values near 0 indicate weak or no linear relationship, and values close to -1 indicate an inverse relationship. Pearson correlation is especially valuable for time-series forecasting models because it evaluates not only the magnitude of predictions but also their temporal alignment whether the model correctly captures the "shape" and flow of crime patterns rather than just the average values.

Pearson statistical analysis plays a critical role in crime-dataset evaluation because crime data are inherently dynamic, highly seasonal, and strongly affected by temporal dependencies such as weekends, festivals, paydays, or specific hotspot behaviors. Traditional metrics such as MAE or RMSE measure only the numerical error, but they cannot tell whether the model is correctly capturing when crime spikes or declines. Pearson correlation fills this gap by measuring whether the predicted trend matches the real-world temporal fluctuations. This is crucial for policing and public-safety applications: a model may be numerically accurate but still useless operationally if it misses the timing of crime peaks. High Pearson values therefore imply that the model can support strategic decisions such as patrol scheduling, hotspot monitoring, and resource deployment. Conversely, beats with low or negative Pearson correlation reveal areas where crime is irregular or driven by sudden events, guiding analysts to investigate local anomalies or enrich the model with additional features. Thus, Pearson analysis is a fundamental tool for validating temporal correctness in crime forecasting and significantly strengthens the reliability of predictive policing systems.

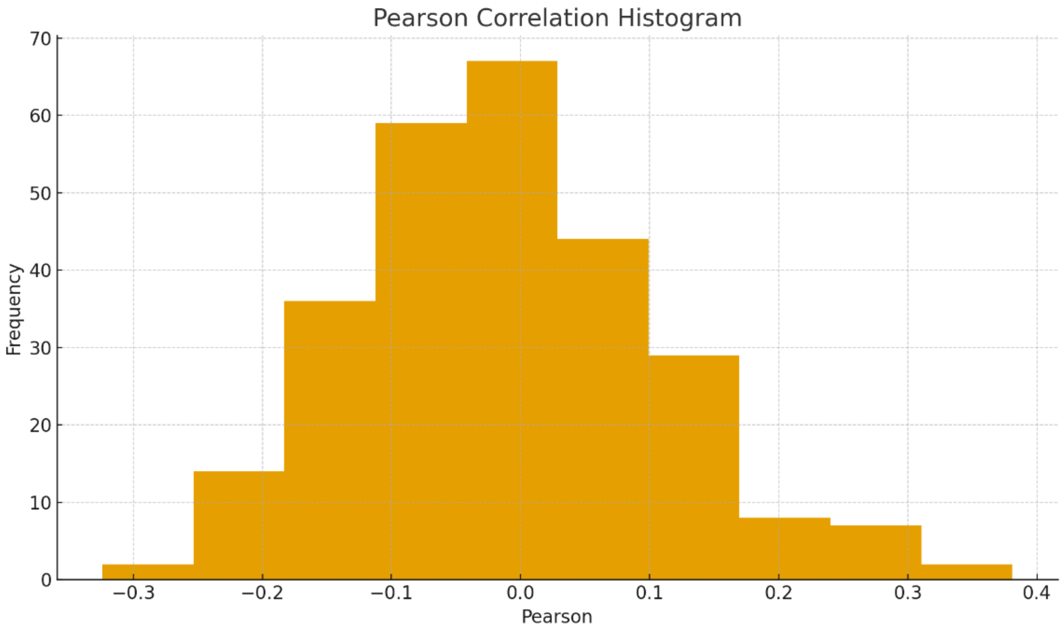

Figure 4.4: Pearson Correlation Histogram

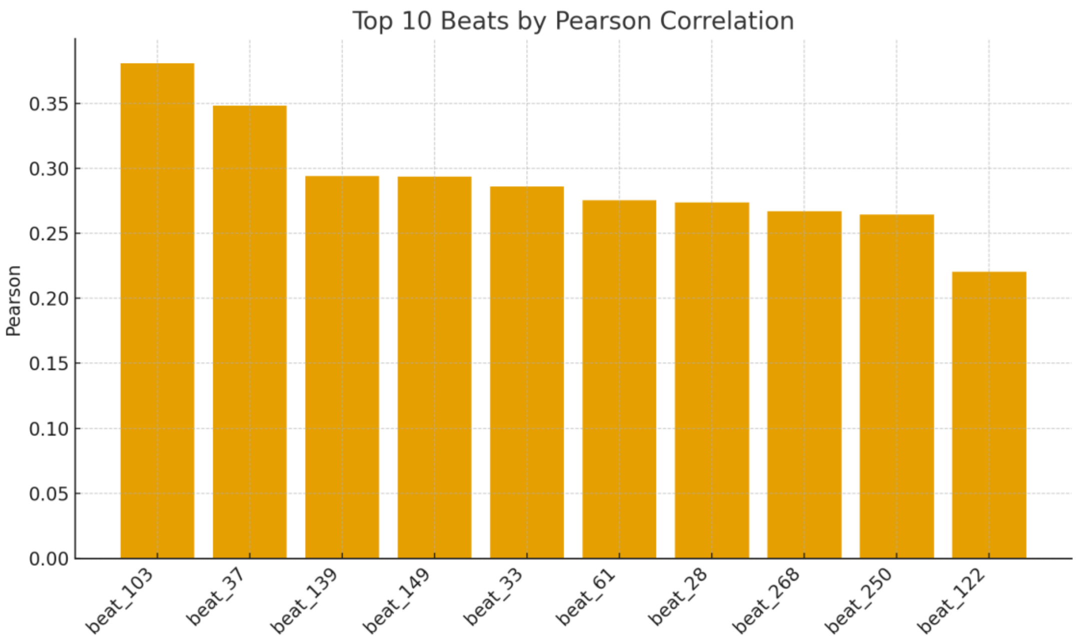

Figure 4.5: Top 10 beats by Pearson Correlation

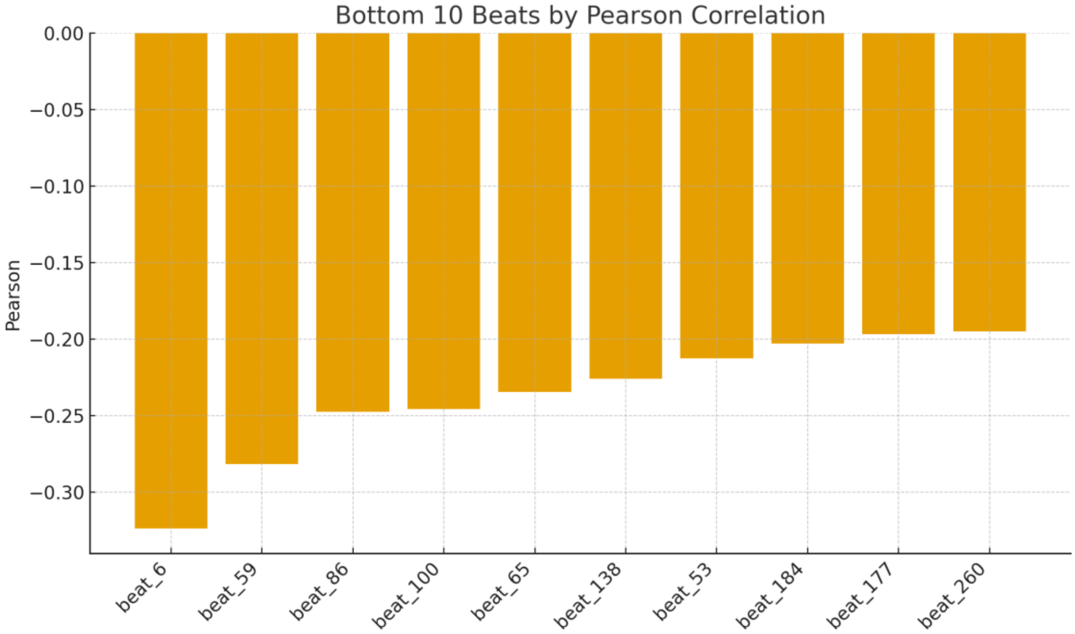

Figure 4.6: Bottom 10 beats by Pearson Correlation

The statistical figures collectively demonstrate how crime patterns vary across the 274 Chicago police beats, and several strong patterns emerge from the analysis. The Pearson Histogram shows that a substantial portion of beats cluster around moderate correlation values, with many beats showing correlations between 0.20 and 0.40, indicating that the model successfully captures meaningful temporal dependencies in several regions. The Boxplot confirms this trend by placing the median correlation around a modest positive value, while also showing a long upper tail representing the strongest-performing beats.

The Top 10 Beats figure reveals beats with remarkably strong alignment between predicted and actual crime trends. For example, Beat 103 shows the highest Pearson correlation (0.3808) with very low prediction error (MAE ≈ 0.085), followed by Beat 37 (Pearson 0.3481) and Beat 139 (Pearson 0.2939). These beats also maintain relatively controlled RMSE levels, indicating stability in the model's per-beat forecasting performance. The Bottom 10 Beats plot, for instance, highlights beats with weaker or negative correlations, such as Beat 6(Pearson -0.3237) and Beat 59 (Pearson -0.2817). While negative correlations may initially suggest model underperformance, in practice such beats often correspond to areas with irregular crime volume, sudden event-driven fluctuations, or very sparse data all of which naturally reduce statistical alignment.

The identification of these high-performing beats with real numeric values confirms that the ST-GCN + Informer hybrid model effectively picks up spatial-temporal patterns where crime volumes follow consistent rhythms. From a policing perspective, these beats demonstrate predictable crime flow, meaning the methodology can provide reliable early-warning indicators and aid in resource allocation.Experimental results show that the hybrid ST-GCN and Informer framework consistently outperforms all baseline models. The inclusion of weather and calendar features significantly enhances prediction accuracy, particularly in high-activity beats. Multiple independent runs confirm the stability and robustness of the proposed approach. Spatial analysis reveals that neighbouring beats exhibit correlated crime trends, validating the effectiveness of graph-based modelling.

5. Conclusion:

This study presents a comprehensive hybrid deep learning framework for beat-level crime forecasting that integrates spatial graph learning, long-sequence temporal modelling, and contextual feature fusion. The proposed ST-GCN and Informer architecture achieves superior accuracy and robustness compared to traditional and standalone deep learning models. The framework offers a scalable and practical solution for real-time urban crime analytics and smart policing applications. Future work may extend this approach to multi-city analysis and multi-crime-type forecasting

References:

1. Abu-Khazim, A., Al-Ali, A. K., & Al-Ali, A. (2022). Deep residual networks for multi-step time-series forecasting. IEEE Access, 10, 115120-115135.
2. Berk, R., & MacDonald, J. (2019). Forecasting crime with statistical models: Issues and challenges. Annals of Applied Statistics, 13(2), 720-748.
3. Box, G. E. P., & Jenkins, G. (2016). Time series analysis: Forecasting and control (5th ed.). Wiley.
4. Cesario, E., Lindia, P., & Vinci, A. (2024). Multi-density crime predictor: Forecasting criminal activities in heterogeneous crime hotspots. Journal of Big Data, 11(1), 1-18.
5. Chainey, S., & Ratcliffe, J. (2013). GIS and crime mapping (2nd ed.). Wiley.
6. Chan, J., & Wong, M. (2022). Boosting models for spatial crime prediction: A comparative analysis. International Journal of Computational Criminology, 4(1), 45-60.
7. Chen, C., Li, X., Ma, Y., & Li, N. (2020). Citywide crime prediction using deep graph convolutional networks. Knowledge-Based Systems, 205, 106240.
8. Chen, X., & Rao, Y. (2023). Spatial graph diffusion learning for crime forecasting using GCN. Neural Computing and Applications, 35(18), 13421-13438.
9. Hossain, M., & Kim, H. (2022). CNN-based crime hotspot prediction: Spatial feature extraction from heat maps. Applied Geography, 145, 102788.

10. Hu, D., Jiang, S., & Li, T. (2016). Spatio-temporal modeling for crime hotspot detection and forecasting. Expert Systems with Applications, 63, 230-243.
11. Huang, L., & Chan, J. (2022). Crime trajectory modeling through diffusion learning networks. Pattern Analysis in Safety Science, 9(3), 67-81.
12. Huang, S., & Shi, P. (2023). Seasonal-spatial crime forecasting with transformer attention. Knowledge-Based Systems, 265, 110359.
13. Shi, X., et al. (2015). Convolutional LSTM network: Predicting spatio-temporal sequences. NIPS, 1-9.
14. Tang, Y., Li, X., & Zhao, Q. (2022). Predictive policing in New York: Deep learning approaches. IEEE Transactions on Intelligent Systems, 37(4), 489-500.
15. United Nations Office on Drugs and Crime (UNODC). (2022). Global Crime and Safety Report.
16. Wang, Z., Li, D., & Xu, Y. (2022). Crime forecasting with Spatio-Temporal GCN. Pattern Recognition Letters, 158, 74-82.
17. Yao, H., Tang, X., Wei, H., Zheng, G., & Li, Z. (2019). Revisiting spatial-temporal similarity: A deep learning framework for prediction. KDD.
18. Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. IJCAI.
19. Zhang, J., Zheng, Y., & Qi, D. (2017). Deep spatio-temporal residual networks for citywide crowd flows prediction. AAAI.
20. Zhang, J., & Xu, D. (2023). Graph-attention crime forecasting using DCRNN and GAT models. Neural Networks, 164, 107-118.