
**3RD WORKSHOP ON
MACHINE LEARNING IN
NETWORKING (MaLeNe)
PROCEEDINGS**

**SEPTEMBER 1,
2025**



**CO-LOCATED WITH
THE 6TH INTERNATIONAL CONFERENCE ON
NETWORKED SYSTEMS (NETSYS 2025)
ILMENAU, GERMANY**

HiSTM: Hierarchical Spatiotemporal Mamba for Cellular Traffic Forecasting

Zineddine Bettouche, Khalid Ali, Andreas Fischer, Andreas Kassler

Deggendorf Institute of Technology

Dieter-Görlitz-Platz 1, 94469 Deggendorf

{zineddine.bettouche, khalid.ali, andreas.fischer, andreas.kassler}@th-deg.de

Abstract—Cellular traffic forecasting is essential for network planning, resource allocation, or load-balancing traffic across cells. However, accurate forecasting is difficult due to intricate spatial and temporal patterns that exist due to the mobility of users. Existing AI-based traffic forecasting models often trade-off accuracy and computational efficiency. We present Hierarchical SpatioTemporal Mamba (HiSTM), which combines a dual spatial encoder with a Mamba-based temporal module and attention mechanism. HiSTM employs selective state space methods to capture spatial and temporal patterns in network traffic. In our evaluation, we use a real-world dataset to compare HiSTM against several baselines, showing a 29.4% MAE improvement over the STN baseline while using 94% fewer parameters. We show that the HiSTM generalizes well across different datasets and improves in accuracy over longer time-horizons.

Index Terms—Time series forecasting, spatiotemporal modeling, 5G network traffic prediction, deep learning, state space models, Mamba architecture, attention mechanisms, convolutional neural networks (CNNs), hierarchical modeling, AI for telecommunications.

I. INTRODUCTION

Accurate traffic forecasting plays a critical role in enabling predictive network resource allocation, network planning and network optimization, directly influencing the operational expenditure of telecom providers [1]. Traditional time series forecasting methods such as ARIMA [2] and AI-based models such as LSTM [3] often treat each base station independently and fail to account for spatial dependencies among neighboring cells. To address this challenge, recent work focused on employing AI-based models jointly with data pre-processing and feature engineering techniques to incorporate spatial knowledge in the training procedure [4], [5].

A key challenge in spatiotemporal modeling is the trade-off between using a single global model and multiple cell-specific models. While cell-specific models can capture local dynamics with higher fidelity, they are costly to train (as they require one model per cell), validate, deploy, and maintain, especially at the scale of modern cellular networks. Conversely, global models may benefit from exploiting shared patterns across cells but often underperform due to distributional heterogeneity [1].

Alternatively, a global model can be trained to cover the entire spatial grid while incorporating the spatial dependencies as features, formulating the forecasting task as a multivariate timeseries forecasting problem. However, such global models often marginally improve performance due to their limited

ability to capture the intricate spatiotemporal dynamics effectively. In 5G networks, user mobility and overlapping coverage areas introduce non-trivial spatial correlations [6], making standalone temporal modeling insufficient as they may introduce large forecasting errors.

Consequently, a single spatiotemporal model capturing both temporal dynamics and latent spatial interactions may exploit correlations among neighboring cells better, leading to increased forecasting accuracy. Indeed, recent advancements using graph neural networks (GNNs) and transformers have shown promise in learning complex spatiotemporal dependencies [7]. However, these architectures typically suffer from high computational cost, making them less viable for real-time or large-scale deployments.

To address these challenges, we propose HiSTM, a novel spatiotemporal forecasting model based on the Mamba architecture [8], leveraging a hierarchical structure that captures both local variability and global patterns. Unlike graph-based or attention-heavy models, HiSTM demonstrates promising performance with reduced resource requirements, offering a favorable trade-off between accuracy and computational efficiency. Our contributions are summarized as follows:

- We propose HiSTM, a novel spatiotemporal forecasting architecture that leverages the standard Mamba model combined with optimized data processing—applying spatial CNNs per frame, temporal modeling over spatial patches via Mamba, and temporal attention to aggregate time-step features for improved representation.
- We apply HiSTM on real-world open-source cellular traffic datasets, demonstrating improved forecasting accuracy (Milan dataset) and consistent generalization ability when evaluated on unseen dataset of a different city (Trentino dataset).
- We show that HiSTM achieves a 29.4% MAE reduction compared to the STN baseline and competitive performance against other methods.
- HiSTM achieves better accuracy over longer forecasting horizons, with a 58% slower error accumulation rate than STN, highlighting its ability to capture long-term temporal dependencies.

The rest of this paper is structured as follows: Section II reviews 5G traffic prediction methods. Section III formalizes the forecasting problem, introduces HiSTM and baselines

we compare against. Section IV describes the datasets and experimental setup. Section V presents results and analysis. Section VI concludes the study and outlines directions for future work.

II. RELATED WORK

Traditionally, 5G traffic forecasting is treated as a purely temporal forecasting task, relying on recurrent neural networks. AI-based models such as vanilla RNNs, LSTMs [3], and GRUs [9] were applied to learn sequential traffic patterns, often achieving gains over classical time-series methods such as ARIMA [2]. However, these approaches inherently ignore spatial correlations. Prior studies note that while RNN-based methods can capture long-term dependencies, they neglect spatial context, reducing their effectiveness [10] [11]. For example, Chen et al. proposed a multivariate LSTM (MuLSTM) using dual streams for traffic and handover sequences, but the model remains fundamentally temporal [12]. Temporal models employ hybrid or preprocessing strategies for improved accuracy. For instance, Hachemi et al. [4] applied an FFT filter before LSTM to separate periodic signals, while Wang and Zhang [5] used Gaussian Process Regression alongside LSTM to manage bursty patterns. These methods often operate at hourly aggregation levels and improve accuracy in long-term forecasts, though they still lack spatial awareness and incur higher computational cost.

Recently, spatial models have become popular as they can exploit spatial correlations within the data. For example, grid-based methods apply ConvLSTM [13] or 3D-CNNs to traffic maps [14], though these require regular cell layouts. Graph-based architectures, such as STCNet and its multi-component variant A-MCSTCNet, combine CNNs or GRUs with attention mechanisms [15]. STGCN-HO uses graph convolutions based on handover-derived adjacency to model spatial links across base stations, with Gated Linear Units (GLUs) for time dependencies [16]. Many of these models supplement deep networks with auxiliary features, such as weekday labels or slice-specific parameters, to improve generalization. More advanced spatiotemporal architectures, such as DSTL, adopt a dual-step transfer learning scheme that clusters gNodeBs and fine-tunes shared RNN models per group [17]. It forecasts at a 10-minute interval while reducing training overhead. To the best of our knowledge, Mehrabian et al. [18] published the first model to incorporate the Mamba framework into a spatiotemporal graph-based predictor for 5G traffic. It adapts a dynamic graph structure within a bidirectional Mamba block to model complex spatial and temporal dependencies.

In contrast, our model, HiSTM, uses the Mamba framework and adapts data input and output processing to work directly on spatial grids with temporal attention layers. This design improves accuracy and efficiency, especially for long-term forecasts and generalization to unseen data.

III. METHODOLOGY

This section formalizes the spatiotemporal forecasting problem, describes the proposed model architectures, and outlines

the baseline models that we compare against.

A. Problem Formulation

Cellular network traffic forecasting aims to predict future traffic volume based on historical observations. We formulate this task as a spatiotemporal sequence prediction problem. Given a sequence of traffic measurements $M = \{M_1, M_2, \dots, M_T\}$ where each $M_t \in \mathbb{R}^{H \times W}$ represents a spatial grid of $H \times W$ cells at time step t , our goal is to predict the next traffic volume grid M_{T+1} . Each grid cell contains a scalar value representing the traffic volume. This volume correlates with resource demands in the cellular network.

We define our input tensor as $X_t^{(i,j)} \in \mathbb{R}^{T \times K \times K}$, where T is the sequence length (window of observation time steps), $K \times K$ represents the spatial dimensions of the input kernel, and (i, j) are the spatial coordinate in the grid M_t . The prediction target is $x_{t+1}^{(i,j)} \in \mathbb{R}$, representing the traffic volume at the center cell of the kernel for the next time step. Formally, we aim to learn a function $f : \mathbb{R}^{T \times K \times K} \rightarrow \mathbb{R}$ that minimizes the prediction error $\min_{\theta} L(f_{\theta}(X), x_{t+1}^{(i,j)})$, where L is the loss function (i.e., Mean Absolute Error), and θ represents the learnable parameters of the model. The spatiotemporal nature of the data introduces a unique challenge: capturing both spatial correlations between neighboring cells and temporal dependencies across the sequence.

B. Proposed Architecture: HiSTM

We propose the Hierarchical SpatioTemporal Mamba (HiSTM), an architecture that combines hierarchical spatiotemporal processing with attention-based temporal aggregation (Figure 1). Given an input tensor $\mathbf{X} \in \mathbb{R}^{T \times K \times K}$ (where T is the number of time steps and $K \times K$ is the spatial kernel), the model predicts target values through three key components:

1) *Hierarchical Spatiotemporal Encoding*: The input $\mathbf{X}$ is first augmented with an initial channel dimension (i.e., $D_{in} = 1$), then passes through N stacked Encoder Layers. Each Encoder Layer l transforms its input $\mathbf{X}^{(l-1)}$ (where $\mathbf{X}^{(0)}$ is the initial augmented input) into an output $\mathbf{X}^{(l)} \in \mathbb{R}^{T \times K \times K \times C}$. The operations within each layer are:

- **Spatial Convolution**: Input features are first reshaped appropriately (e.g., to $T \times D'_{in} \times K \times K$). A 2D convolution followed by a ReLU activation is applied. The first layer up-projects D'_{in} to C channels. Subsequent layers take $D'_{in} = C$ channels and output C channels.
- **Temporal Mamba Processing**: The C -channel output from the convolution is reshaped to $\mathbf{X}_{\text{flat}} \in \mathbb{R}^{(K^2) \times T \times C}$. A Mamba SSM [8] then models the temporal dependencies for each of the K^2 spatial locations treated as sequences of length T , with $d_{\text{mamba}} = C$.

We use Mamba for its state space foundation, which models sequences through continuous-time dynamics rather than attention. This enables precise control over temporal structure and inductive bias, making it suitable for forecasting tasks where long-range temporal dependencies interact with fine-grained spatial patterns. Its ability to selectively retain and




Figure 1. HiSTM Architecture

propagate information aligns with the demands of spatiotemporal modeling.

The Mamba output is reshaped back to $\mathbb{R}^{T \times K \times K \times C}$, forming $\mathbf{X}^{(l)}$. The output of this encoding stage is $\mathbf{X}_{\text{encoded}} = \mathbf{X}^{(N)}$.

2) *Temporal Attention-Based Aggregation*: From the encoded features $\mathbf{X}_{\text{encoded}} \in \mathbb{R}^{T \times K \times K \times C}$, features corresponding to the center spatial cell are extracted across all T time steps. This yields a sequence $\mathbf{X}_{\text{center}} \in \mathbb{R}^{T \times C}$. An attention mechanism then computes an aggregated context vector $\mathbf{c} \in \mathbb{R}^{B \times C}$:

$$e_t = \text{Linear}_{\text{att}}(\mathbf{h}_t), \quad \alpha_t = \frac{\exp(e_t)}{\sum_{j=1}^T \exp(e_j)}, \quad \mathbf{c} = \sum_{t=1}^T \alpha_t \mathbf{h}_t \quad (1)$$

where $\mathbf{h}_t \in \mathbb{R}^C$ is the feature vector from $\mathbf{X}_{\text{center}}$ (for a given batch instance) at time step t . The $\text{Linear}_{\text{att}}$ layer maps from $\mathbb{R}^C \rightarrow \mathbb{R}$, producing a scalar energy e_t . The softmax function normalizes these energies across all T time steps to obtain attention weights α_t .

3) *Prediction Head*: Finally, the aggregated context vector $\mathbf{c} \in \mathbb{R}^C$ is fed into a Multilayer Perceptron (MLP) head. This MLP consists of two linear layers with a ReLU activation in between, passing the dimensionality from C to MLP_{in} , and then to 1, to produce the final prediction $\hat{y} \in \mathbb{R}^1$.

C. Baseline Models

We benchmark HiSTM against the following baselines:

- **STN** [14]: A deep neural network capturing spatiotemporal correlations for long-term mobile traffic forecasting.
- **xLSTM** [19]: A scalable LSTM variant with exponential gating and novel memory structures (sLSTM/mLSTM) designed as Transformer alternatives. For fair comparison, our implementation uses only one mLSTM layer (denoted as xLSTM[1:0] in their paper) to prioritize computational efficiency while retaining its parallelizable architecture.
- **STTRE** [20]: A Transformer-based architecture leveraging relative embeddings to model dependencies in multivariate time series.




Figure 2. Spatial distribution of traffic flow intensity across Milan and Trentino regions.

- **VMRNN-B & VMRNN-D** [21]: Vision Mamba-LSTM hybrids addressing CNN's limited receptive fields and ViT's computational costs. VMRNN-B (basic) and VMRNN-D (deep) use Mamba's selective state-space mechanisms for compact yet competitive performance.

IV. DATASETS & EXPERIMENTAL SETUP

This section outlines the datasets, preprocessing steps, training configuration, and evaluation process used to assess model performance in spatiotemporal data forecasting.

A. Datasets

We use the dataset introduced by Barlacchi et al. [22], which contains two sub-datasets from Milan and Trentino with 10,000 and 11,466 spatial cells, respectively. Both datasets record SMS-in/out, Call-in/out, and Internet Traffic Activity in 10-minute intervals. The data exhibits spatial heterogeneity (see Figure 2), with higher activity in central regions and localized clusters elsewhere, adding complexity to forecasting tasks. To better visualize these differences, the Milan heatmap uses a linear scale, while the Trentino heatmap employs a logarithmic scale to highlight the dense urban center against its sparsely populated surroundings.

Lag plot analysis (Figure 3) reveals that both datasets maintain correlation at higher lag values. Individual cell traffic demonstrates higher volatility (Approximate Entropy [23]) of




Figure 3. Lag plot for the autocorrelation of the dataset for the entire grid.

1.386 for a single cell) compared to spatially aggregated traffic (0.196), with the latter exhibiting enhanced cyclical patterns and passing the Augmented Dickey Fuller test [24] for stationarity. It is evident that the aggregated series is more correlated with itself for different lags; hence, it is easier to predict. Consequently, incorporating the spatial element in the prediction can help the model to capture the distribution more effectively, reduce the influence of individual events and magnify predictable cyclical patterns. This makes the series more predictable in a spatiotemporal context compared to when using individual per-cell series.

B. Data Preprocessing

To ensure robust spatiotemporal feature extraction, chronological integrity, and leakage-free normalization for model training and evaluation, we preprocess the raw 100×100 spatial grid as follows. We set $K = 11$, with boundary effects mitigated by cropping. Temporal sequences are constructed by aggregating six consecutive time steps. For the training set, input sequences are generated using a temporal stride of 6 between consecutive samples, minimizing temporal overlap from a single grid to promote feature diversity. During testing, a stride of 1 is used to ensure exhaustive evaluation across all temporal segments. The model is trained to predict the central value of the 11×11 kernel at the seventh (subsequent) time step. The preprocessed dataset is chronologically partitioned into training (70%), validation (15%), and test (15%) sets to preserve temporal order and prevent information leakage from future to past. Input features and target values are normalized to the $[0, 1]$ range via Min-Max scaling, where scaling parameters are derived from the training data. Validation and test sets are transformed using these parameters, and out-of-range values are clipped to maintain the bounds.

C. Implementation and Training Configuration

We implement the HiSTM model in PyTorch, leveraging GPU-optimized operations for its Mamba and Conv2D modules to ensure computational efficiency. For reproducibility,

we release the complete source code for our HiSTM implementation in a public repository [25]. For training and inference, we use an AI-server with a single NVIDIA A100 80GB GPU with 64 CPU cores and 512 GB RAM, using CUDA 12.4 and PyTorch 2.6.0+cu124. HiSTM and baseline models were trained with a batch size of 128 for up to 40 epochs, using early stopping with a patience of 15 and saving the best model based on validation loss. We use Adam optimizer with a learning rate of 10^{-4} and ReduceLROnPlateau scheduler (patience: 7, factor: 0.5).

D. Evaluation Metrics

We evaluate prediction accuracy using four complementary metrics. Mean Absolute Error (MAE) measures the average absolute difference between predictions and ground truth. Root Mean Squared Error (RMSE) penalizes large errors more heavily. The coefficient of determination (R^2) quantifies the proportion of variance explained by the model. Structural Similarity Index (SSIM) assesses the visual quality of spatial predictions. All metrics are computed after reversing the normalization to original scale.

V. RESULTS AND ANALYSIS

This section evaluates forecast accuracy, cross-dataset generalization, and computational efficiency, comparing our proposed HiSTM model against baselines.

A. Prediction Accuracy

1) *Single-step Prediction Results:* HiSTM achieves the best single-step forecasting performance on the Milan dataset, with an MAE of 5.2196 and SSIM of 0.9925 (see Table I). Numbers are averaged across all spatial cells. This corresponds to a 29.4% MAE reduction over STN and a 2.3% SSIM gain, while outperforming other architectures such as STTRE (MAE: 5.5558) and parameter-intensive models like VMRNN-D (MAE: 6.4151). HiSTM also reports the lowest RMSE (11.2476) and the highest R^2 score (0.9799), indicating both lower large-error impact and improved variance explanation.

In addition to numerical metrics, Figure 4 provides a density-based comparison of predicted vs. actual values for all models. HiSTM shows a tighter concentration of points along the diagonal, particularly at higher traffic volumes, indicating improved predictive fidelity across the full dynamic range. It also exhibits fewer high-error outliers and lower dispersion compared to other models. STN and VMRNN-B, by contrast, display broader scatter, especially in the high-traffic regime, with more frequent large deviations. The distributional compactness in HiSTM suggests better generalization and robustness across heterogeneous spatial patterns.

2) *Multi-step Autoregressive Forecasting:* HiSTM demonstrates improved stability over extended forecasting horizons (Table II). While all models accumulate error with each additional step, HiSTM maintains the lowest MAE and RMSE across all six steps. At step 6, HiSTM reports an MAE of 6.69, which is 36.8% lower than STN (10.59) and 11.3% below STTRE (7.54). The slope of MAE progression is 58% lower




Figure 4. Density comparison of predicted vs. actual values over the entire grid for HiSTM and baseline models, with associated MAE

Table I
SINGLE-STEP PREDICTION PERFORMANCE ON MILAN DATASET. BEST MODEL INDICATED THROUGH BOLD FONT.

Model	MAE ↓	RMSE ↓	R ² Score ↑	SSIM ↑
STN	7.3908	16.8824	0.9546	0.9853
VMRNN-B	7.1659	19.0876	0.9420	0.9843
xLSTM	6.4672	15.0901	0.9637	0.9870
VMRNN-D	6.4151	16.3284	0.9575	0.9873
STTRE	5.5558	11.4426	0.9791	0.9917
HiSTM	5.2196	11.2476	0.9799	0.9925

than that of STN, indicating reduced error propagation. SSIM also degrades more gradually in HiSTM, remaining above 0.95 even at step 6. This suggests that HiSTM better preserves temporal dependencies and structural consistency across time.

3) *Cross-dataset Generalization on Trentino dataset:* On the unseen Trentino dataset, HiSTM achieves a 47.3% MAE reduction (1.3870 vs. 2.6344) and a 36.9% RMSE improvement (4.8134 vs. 7.6370) over STN (see Table III). It also yields the highest SSIM (0.9916) and R^2 score (0.9649), confirming strong structural preservation. HiSTM outperforms all baselines across all evaluation metrics, including STTRE (MAE: 1.8132), VMRNN-D (MAE: 1.5870), and VMRNN-B (MAE: 1.9751). This demonstrates HiSTM’s capacity to generalize to new spatial environments with distinct activity distributions and scales.

B. Cell-specific Modeling and Spatially-aware Accuracy

To evaluate the spatially-aware accuracy of the model, we use Milan dataset and select a random 7-day time period (1008 time steps) from the test dataset. We use a 6-step memory window and an 11-step kernel on a 100×100 sensor grid to forecast a single next timestep. To analyze spatial performance across varying traffic conditions, we select four representative cells, corresponding to (a) urban (b) suburban (c) rural and (d) the cell with the maximum temporal variance (Figure 5). The predicted versus actual traffic trajectories for all four cells are shown in Figure 6.

The model achieved its lowest MAPE (8.49%) for the urban cell, demonstrating strong performance in dense traffic zones. Prediction accuracy remained stable for the suburban cell (11.62% MAPE) with only minor deviations during peak periods. The rural cell showed higher relative error (24.86%




Figure 5. Selected cells from the Milan’s traffic network. The cells represent different traffic patterns: (a) urban, (b) suburban, (c) rural, and (d) maximum variance cell.

MAPE), where low absolute traffic volumes magnified percentage errors. Notably, the model maintained robust performance (15.30% MAPE) for the high-variance cell despite its extreme fluctuations, highlighting HiSTM’s capacity to handle volatile temporal patterns while showing expected limitations in sparse rural conditions where signal-to-noise ratios are challenging. These performance differences reflect the interaction between model complexity and regional traffic characteristics: urban and suburban areas offer rich temporal patterns and denser signals that align well with the model’s spatiotemporal structure, whereas rural areas pose challenges due to sparse activity and fewer recurring patterns. The high-variance cell underscores the model’s ability to generalize to non-stationary behavior, though peak underestimation suggests further headroom for improvement in handling extremes.

C. Computational Efficiency Analysis

HiSTM achieves a better performance while requiring fewer model parameters (Table IV). At 33.8K parameters (0.13 MB), it requires 18× fewer parameters than VMRNN-D and 5.4× less than xLSTM while delivering faster inference (1.19 ms) than all baselines, except xLSTM. Though its multiply-accumulate operations (MACs, a standard metric for computational workload) total 1.36×10^7 —higher than

Table II

PERFORMANCE COMPARISON OF MODELS OVER MULTIPLE STEPS (AUTOREGRESSIVE FORECASTING). BEST MODEL INDICATED THROUGH BOLD FONT.

Step	HiSTM			STN			STTRE			xLSTM			VMRNN-B			VMRNN-D		
	MAE	RMSE	SSIM	MAE	RMSE	SSIM	MAE	RMSE	SSIM	MAE	RMSE	SSIM	MAE	RMSE	SSIM	MAE	RMSE	SSIM
1	3.87	9.54	0.9833	5.28	13.05	0.9742	4.21	9.83	0.9808	4.64	11.86	0.9759	5.11	13.75	0.9745	4.62	12.25	0.9777
2	4.38	10.95	0.9785	6.44	16.29	0.9615	4.88	11.45	0.9739	5.62	14.63	0.9639	5.08	13.85	0.9740	4.61	12.36	0.9772
3	4.85	12.06	0.9735	7.39	18.54	0.9490	5.47	12.73	0.9668	6.47	16.78	0.9509	6.06	16.33	0.9637	5.35	14.30	0.9698
4	5.56	13.42	0.9680	8.59	21.05	0.9335	6.27	14.12	0.9586	7.50	18.95	0.9356	6.83	17.64	0.9561	6.08	15.49	0.9641
5	6.09	14.62	0.9633	9.55	22.94	0.9196	6.88	15.28	0.9513	8.34	20.58	0.9217	7.45	18.91	0.9492	6.59	16.54	0.9591
6	6.69	16.02	0.9578	10.59	25.01	0.9033	7.54	16.60	0.9430	9.23	22.38	0.9064	8.24	20.38	0.9398	7.30	17.87	0.9520

Table III

SINGLE-STEP GENERALIZATION PERFORMANCE ON TRENTINO DATASET. BEST MODEL INDICATED THROUGH BOLD FONT.

Model	MAE ↓	RMSE ↓	R ² Score ↑	SSIM ↑
STN	2.6344	7.6370	0.9116	0.9762
xLSTM	2.5974	8.9235	0.8793	0.9615
VMRNN-B	1.9751	6.6270	0.9334	0.9839
STTRE	1.8132	5.0050	0.9620	0.9903
VMRNN-D	1.5870	5.5754	0.9529	0.9885
HiSTM	1.3870	4.8134	0.9649	0.9916

Table IV
MODEL COMPARISON

Model	Parameter Count	Size (MB)	GPU (MB)	Inference (ms)	MACs
xLSTM	607,753	2.32	13.69	1.01	5.96×10^5
VMRNN-B	137,282	0.52	9.77	8.16	2.06×10^7
VMRNN-D	1,506,498	5.75	15.47	18.58	5.52×10^7
STTRE	165,380	0.63	58.07	4.54	2.83×10^7
STN	576,755	2.20	11.34	2.46	2.31×10^6
HiSTM	33,794	0.13	10.63	1.19	1.36×10^7

STN’s—HiSTM balances computational cost with accuracy, achieving a MAC/MAE ratio 2.6× better than STTRE’s. This efficiency-profile positions HiSTM as a practical solution for resource-constrained deployment scenarios.

VI. CONCLUSION

In this paper, we presented HiSTM, a hierarchical spatiotemporal model for efficient and accurate cellular traffic forecasting. By combining dual spatial encoders with a Mamba-based temporal module and an attention mechanism, HiSTM captures complex spatiotemporal patterns with minimal overhead. Experiments on real-world datasets show that HiSTM demonstrates competitive performance, reducing MAE by 29.4% compared to STN baseline while maintaining computational efficiency. It generalizes well to the unseen Trentino dataset and sustains lower errors over longer forecast horizons. These results suggest HiSTM’s potential as an efficient approach for cellular traffic prediction, though broader validation across diverse network conditions would strengthen deployment recommendations.

Future work will explore mixture-of-experts approaches to better model spatially-clustered traffic patterns and kernel-to-kernel forecasting to capture finer-grained temporal dynamics




Figure 6. Week-long traffic predictions from the HiSTM model across four representative cells in Milan’s traffic network. The cells were selected to represent different traffic patterns: (a) urban, (b) suburban, (c) rural, and (d) maximum variance cell. Time is shown in days, with each day containing 144 readings (10-minute intervals)

across service types. We aim to investigate diffusion-based decoding strategies to enhance long-range predictive capabilities and extend the model’s scope through different aggregation strategies.

To strengthen generalizability, we will evaluate HiSTM on diverse geographical datasets beyond Italy and implement attention visualization for improved interpretability. We also plan to address practical deployment challenges including

missing data handling, concept drift adaptation, and spatial heterogeneity through adaptive weighting mechanisms.

ACKNOWLEDGEMENT

This work was partly funded by the Bavarian Government by the Ministry of Science and Art through the HighTech Agenda (HTA).

REFERENCES

- [1] W. Shao, J. Wang, Z. Guo, M. Xu, and Y. Wang, "Spatial-temporal neural network for wireless network traffic prediction," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 3, pp. 2104–2113, 2020.
- [2] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Hoboken, New Jersey: Wiley, 2015.
- [3] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [4] H. Hachemi, V. Vidal, and K. Boussetta, "Towards real-time mobile traffic prediction with fft-lstm," in *2020 IEEE GLOBECOM*, 2020.
- [5] J. Wang and T. Zhang, "Gaussian process assisted lstm for 5g traffic prediction," *IEEE Transactions on Network and Service Management*, vol. 17, no. 4, pp. 2472–2485, 2020.
- [6] B. Yu, H. Yin, Z. Zhu, and Q. Zhang, "Spatiotemporal graph neural network for urban traffic flow prediction," *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [7] R. Gu, Q. Liu, X. Li, and Q. Zhu, "Glsttn: A global-local spatial-temporal transformer network for traffic prediction," *IEEE Access*, vol. 9, pp. 152 323–152 334, 2021.
- [8] A. Gu, T. Dao, A. Rudra, B. Recht, and T. B. Hashimoto, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2024.
- [9] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734.
- [10] X. Ma, J. Zhang, and X. S. Shen, "Traffic prediction for mobile networks using machine learning: A survey," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2141–2169, 2019.
- [11] Y. Li, L. Duan, H. Liu, and X. Wang, "A survey on deep learning techniques in wireless resource allocation for 5g and beyond," *IEEE Wireless Communications*, vol. 28, no. 5, pp. 152–159, 2021.
- [12] L. Chen, T.-M.-T. Nguyen, D. Yang, M. Nogueira, C. Wang, and D. Zhang, "Data-driven c-ran optimization exploiting traffic and mobility dynamics of mobile users," *IEEE Transactions on Mobile Computing*, vol. 20, no. 5, pp. 1773–1788, 2021.
- [13] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [14] C. Zhang and P. Patras, "Long-term mobile traffic forecasting using deep spatio-temporal neural networks," in *Proceedings of the 25th Annual International Conference on Mobile Computing and Networking*, 2019, pp. 1–15.
- [15] I. Narmanlioglu and A. Cicek, "Multi-component spatio-temporal modeling for cellular traffic forecasting," *IEEE Access*, vol. 10, pp. 12 771–12 784, 2022.
- [16] C. Borcea, J.-M. Gorce, and M. Dufloy, "Stgcn-ho: Handover-aware spatiotemporal graph convolutional network for mobile traffic forecasting," *IEEE Transactions on Mobile Computing*, 2023.
- [17] S. Aziz, L. Hu, and Z. Han, "Dstl: Dual-step transfer learning for spatiotemporal 5g traffic forecasting," *IEEE Transactions on Network and Service Management*, 2025, to appear.
- [18] S. Mehrabian, L. Jiang, and S. Lee, "A-gamba: Adaptive graph mamba network for 5g spatiotemporal traffic forecasting," *IEEE Transactions on Mobile Computing*, 2025, to appear.
- [19] M. Beck, K. Pöppel, M. Spanring, A. Auer, O. Prudnikova, M. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter, "xlstm: Extended long short-term memory," 2024. [Online]. Available: <https://arxiv.org/abs/2405.04517>
- [20] A. Deihim, E. Alonso, and D. Apostolopoulou, "Sttre: A spatio-temporal transformer with relative embeddings for multivariate time series forecasting," *Neural Networks*, vol. 168, pp. 549–559, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0893608023005361>
- [21] Y. Tang, P. Dong, Z. Tang, X. Chu, and J. Liang, "Vmrnn: Integrating vision mamba and lstm for efficient and accurate spatiotemporal forecasting," 2024. [Online]. Available: <https://arxiv.org/abs/2403.16536>
- [22] G. Barlacchi, M. D. Nadai, R. Larcher, A. Casella, C. Chitic, G. Torrisi, F. Antonelli, A. Vespignani, A. Pentland, and B. Lepri, "A multi-source dataset of urban life in the city of milan and the province of trentino," *Scientific Data*, vol. 2, p. 150055, 2015. [Online]. Available: <https://www.nature.com/articles/sdata201555>
- [23] S. M. Pincus, "Approximate entropy as a measure of system complexity," *Proceedings of the National Academy of Sciences*, vol. 88, no. 6, pp. 2297–2301, 1991.
- [24] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, vol. 74, no. 366a, pp. 427–431, 1979.
- [25] Z. Bettouche and K. Ali, "Hstm: Hierarchical spatiotemporal mamba," <https://github.com/ZineddineBtc/HiSTM-Hierarchical-Spatiotemporal-Mamba>, 2025.