

# Data-driven Modeling of Combined Sewer Systems for Urban Sustainability: An Empirical Evaluation

Vipin Singh<sup>1,\*</sup>, Tianheng Ling<sup>2</sup>, Teodor Chiaburu<sup>1</sup> and Felix Biessmann<sup>1,3</sup>

<sup>1</sup>Berlin University of Applied Sciences and Technology, Luxemburger Str. 10, 13353 Berlin, Germany

<sup>2</sup>University of Duisburg-Essen, Bismarckstraße 90, 47057 Duisburg, Germany

<sup>3</sup>Einstein Center for Digital Future, Wilhelmstraße 67, 10117 Berlin, Germany

## Abstract

Climate change poses complex challenges, with extreme weather events becoming increasingly frequent and difficult to model. Examples include the dynamics of Combined Sewer Systems (CSS). Overburdened CSS during heavy rainfall will overflow untreated wastewater into surface water bodies. Classical approaches to modeling the impact of extreme rainfall events rely on physical simulations, which are particularly challenging to create for large urban infrastructures. Deep Learning (DL) models offer a cost-effective alternative for modeling the complex dynamics of sewer systems. In this study, we present a comprehensive empirical evaluation of several state-of-the-art DL time series models for predicting sewer system dynamics in a large urban infrastructure, utilizing three years of measurement data. We especially investigate the potential of DL models to maintain predictive precision during network outages by comparing global models, which have access to all variables within the sewer system, and local models, which are limited to data from a restricted set of local sensors. Our findings demonstrate that DL models can accurately predict the dynamics of sewer system load, even under network outage conditions. These results suggest that DL models can effectively aid in balancing the load redistribution in CSS, thereby enhancing the sustainability and resilience of urban infrastructures.

## Keywords

Urban Sustainability, Combined Sewer Overflow, Deep Learning, Time-series Forecasting

## 1. Introduction and Related Work

Climate change has increased the frequency and intensity of extreme weather events [1], which pose significant challenges to urban infrastructure and environmental management [2]. Managing Combined Sewer Systems (CSS) [3]. Heavy rainfall can overwhelm the capacity of these systems, leading to overflows that release untreated sewage into rivers and lakes [4]. This contamination compromises water quality and poses direct risks to human health [5].

Many urban areas that utilize CSS have implemented overflow basins to mitigate this risk [4], as shown in Figure 1. However, there remains a significant gap in understanding the dynamics of water levels in these overflow basins. Traditional methods for modeling the dynamics of sewer systems rely on physical simulations [6]. These systems are challenging to apply to large urban infrastructures as they require domain expertise and detailed data on the system components, which is often unavailable or imposes significant financial costs. Improving the forecasting of these water levels can significantly

enhance real-time flow control and inform maintenance and extension planning for sewage overflows.

Data-driven approaches, such as Deep Learning (DL) models, particularly time-series models, offer a promising alternative for modeling sewer system dynamics. Any target variable can be modeled with a combination of variables of the sewer system and exogenous variables, such as rainfall, without an explicit model of the sewer system, as it would be required for classical hydrological systems. These models enable accurate and flexible modeling of sewage treatment facilities to proactively manage and redistribute the load, thus preventing overflows and mitigating their impacts [7]. Such predictive capabilities are crucial for timely interventions and informed decision-making in urban water management.


Transforming urban infrastructure to address the challenges of extreme weather on urban sewer systems requires cost-efficient modeling. Here, we investigate the potential of data-driven methods and demonstrate the potential of DL-based approaches to foster a more equitable and sustainable urban environment. Our data-driven models enhance the predictive capabilities of sewer system management, reducing the likelihood of untreated sewage overflows while keeping financial costs lower than traditional methods, thereby making proactive management accessible to lower-income and marginalized communities. These models also protect public health by minimizing water contamination, particularly in vulnerable populations, and ensure cleaner water bodies.

*2nd Workshop on 'Public Interest AI' co-located with the 47th German Conference on AI (KI 2024), September 23, 2024, Julius-Maximilians Universität Würzburg, Germany.*

\*Corresponding author.

✉ vipin.singh@bht-berlin.de (V. Singh); tianheng.ling@uni-due.de (T. Ling); chiaburu.teodor@bht-berlin.de (T. Chiaburu); felix.biessmann@bht-berlin.de (F. Biessmann)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



**Figure 1:** CSS with overflow basins collect rainwater and wastewater into multiple basins. During heavy rainfall, these basins may exceed capacity, leading to the untreated mixture overflow into the environment [10].

Additionally, this solution enables cities to optimize resource allocation and emergency responses, reducing environmental damage and long-term economic costs. It underscores the importance of deploying advanced technologies in ways that prioritize the well-being of all community members, ensuring that technological advancements contribute to the common good and support populations in the face of climate change, aligning with the principles of public interest AI [8].

It is worth mentioning that there are approaches aiming at combining classical physics-based simulations with the flexibility of DL methods [9]. The challenge with these combinations is that the model architectures require the same detailed knowledge about the sewer system as traditional methods. In addition, modeling the mixed viscosity of wastewater imposes significant complexity on the physical models. Here we focus on a data-driven approach that learns the system dynamics from measurements and is thus readily applicable without building cost-intensive digital twins reflecting the physical properties of the sewer system.

In this work, we evaluate the performance of various advanced time series models for forecasting the water levels in CSS overflow basins. We explore and compare two approaches to time series forecasting: (1) a global model approach, which incorporates a number of exogenous variables, such as rainfall data, and (2) a local model approach, which relies solely on historical water level data. Our analysis aims to identify the most effective models and approaches for practical application in sewage overflow management. Our study extends prior work on time series models for wastewater modeling [10]

with the following main contributions:

- **Comprehensive Model Evaluation:** We systematically evaluated multiple state-of-the-art time series models based on three years of real-world data. The Long Short-Term Memory (LSTM) [11] and Temporal Fusion Transformer (TFT) [12] models, in particular, showed superior performance, with LSTM achieving the lowest Mean Squared Error (MSE) and TFT providing robust and consistent predictions across various conditions.
- **Global vs. Local Model Comparison:** We compared global and local model approaches, highlighting their strengths and limitations in sewage overflow forecasting. Our findings indicate that global models generally outperform local models in terms of MSE. However, local models are advantageous in scenarios where exogenous data is unavailable, offering a computationally efficient alternative.

## 2. Data and Preprocessing

In our study, we utilized data<sup>1</sup> provided by *Wirtschaftsbetriebe Duisburg*<sup>2</sup>. The dataset comprises time series sensory data, including water levels in rain basins and water tanks, energy consumption of pumps, and rainfall

<sup>1</sup>Data cannot be made publicly available. Readers can contact the corresponding author for details.

<sup>2</sup><https://www.wb-duisburg.de>

amounts. The sensor data were collected from six locations in Duisburg’s Vierlinden district, covering three years, from January 1, 2021, at 00:00 AM until January 1, 2024, at 00:00 AM. The recording intervals are irregular as the sensors were read out event-based, with sensor update intervals ranging from 1 second to 1 hour.

To standardize the data, we resampled it by calculating the mean values closest to each full-hour mark. This resampling procedure resulted in a total of 26,280 data points, with each data point comprising 35 features derived from the sensory data across the different locations. For missing values in the rainfall measurements, we utilized data from the nearest weather station of the *Deutsche Wetterdienst*<sup>3</sup>, specifically the station in Duisburg-Baerl, located 4.5 km from the sewage treatment plant that recorded the rainfall. For the other features, linear interpolation was employed to fill the missing values. Additionally, an indicator column was added for each feature with missing values to denote whether the corresponding value was interpolated.

### 3. Methodology

This section details the time series models employed in this study, including the selection, implementation, and comparisons of global and local model approaches.

#### 3.1. Neural Network Architectures

For our empirical evaluation, we selected six state-of-the-art neural time series models. While classical regression models, such as tree-based methods, can be effective for time series data, the state of the art in water modeling increasingly relies on neural network models [13]. We thus focus on these models based on their effectiveness and versatility in forecasting tasks. The selected models are:

- LSTM [11]: LSTM networks are well-suited for capturing long-term dependencies in sequential data, making them ideal for time series forecasting.
- DeepAR [14]: This probabilistic forecasting model leverages autoregressive recurrent networks, providing robust predictions with uncertainty estimates.
- Neural Hierarchical Interpolation for Time Series Forecasting (N-HiTS) [15]: As a neural hierarchical time series model, N-HiTS excels in capturing complex temporal patterns.

- Transformer [16]: Originally designed for nature language processing [17], Transformers use attention mechanism to effectively capture relationships across different time steps [16].
- Temporal Convolutional Network (TCN) [18]: TCNs can model long-range dependencies in time series data while being computationally efficient.
- TFT [12]: TFT combines the strengths of LSTM and attention mechanisms to provide interpretable and accurate forecasts.

These models were implemented using the *Darts*<sup>4</sup> Python library, which offers a user-friendly interface for time-series forecasting. *PyTorch*<sup>5</sup> was used as the supporting framework.

#### 3.2. Global vs Local Model Approach

Our study considers two approaches to time series forecasting, each motivated by distinct real-world scenarios.

1. **Global Model Approach:** This approach corresponds to the scenario of normal CSS operation, where all sensors are fully operational, and all data can be transmitted reliably over the network. In this case, all available data, including exogenous variables such as rainfall data, are integrated into a single model for forecasting the relevant target variable. This approach, referred to as the *global model*, allows the models to leverage additional contextual information to improve forecasting precision.
2. **Local Model Approach:** This approach is designed for scenarios where not all data is available. Such circumstances can arise when sensors are damaged or network connections are unstable due to extreme weather events, or security incidents. To mimic these cases, predictions are made using only the historical data from the specific sensor in question, without access to additional contextual information. This approach is referred to as the *local model*. It is intended for future deployment on edge devices, enabling localized and resilient forecasting capabilities [19].

Overall, the global model approach leverages extensive data to enhance forecasting precision, whereas the local model approach ensures robustness and adaptability in environments with limited data availability. By evaluating both approaches, we aim to provide a comprehensive solution for diverse operational scenarios in urban wastewater management.

<sup>3</sup>[https://www.dwd.de/DE/Home/home\\_node.html](https://www.dwd.de/DE/Home/home_node.html)

<sup>4</sup><https://unit8co.github.io/darts/>

<sup>5</sup><https://pytorch.org>

## 4. Experiments

This section introduces the experimental settings, including data splits, model development, and error metrics.

### 4.1. Datasets

The dataset was divided into training, validation, and testing sets. The first two years of data were used for training and validation, while the last year was reserved for testing. Within the initial two years, 80% of the data was allocated for training and 20% for validation. Given the sequential nature of time series data, the data were not shuffled, and the split was performed in chronological order to maintain temporal dependencies. To ensure consistency across features, standard scaling was applied using default parameters ( $\mu = 0$ ,  $\sigma^2 = 1$ ) prior to model training.

After fine-tuning (see Appendix A), we determined that a 72-hour input sequence was optimal for forecasting a 12-hour prediction sequence. This prediction sequence was determined together with domain experts to meet operational requirements.

### 4.2. Model Development

To prevent overfitting, early stopping with patience of 10 epochs was adopted. All models were optimized using the *Adam* optimizer [20] with 32-bit floating point precision. Training sessions were conducted on an NVIDIA A100 with 40GB VRAM, utilizing CUDA 12.2 and Python 3.10. Hyperparameter optimization was performed using the Tree-structured Parzen Estimator algorithm provided by the *Optuna* library<sup>6</sup>. Each model had a training budget of 600 trials, with each trial consisting of 100 epochs. Further details on the hyperparameter optimization process are provided in Appendix A. The best hyperparameter configuration for each model was then evaluated using 100 different random weight initializations to ensure robust comparisons.

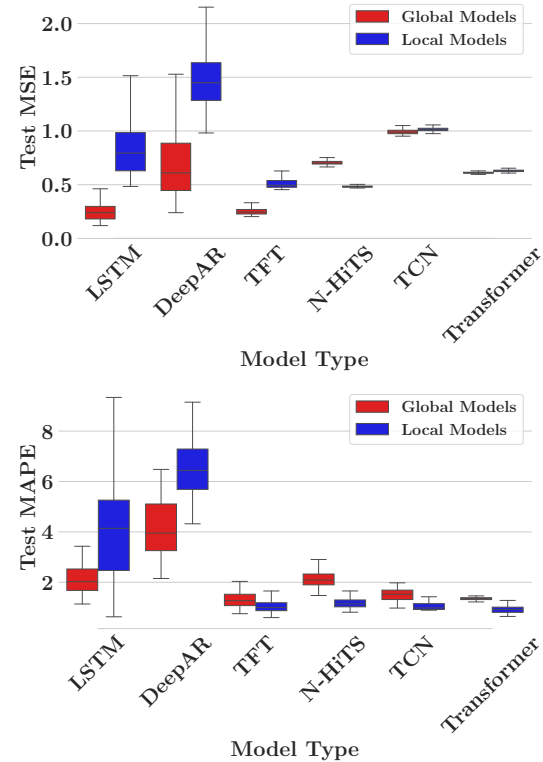
### 4.3. Error Metrics

We compared various error metrics well established in the field of time series forecasting [21], including MSE and Mean Absolute Percentage Error (MAPE). For model training, MSE was selected as the loss function, except for the probabilistic model DeepAR, which utilized the negative log-likelihood. While both metrics have limitations, they are among the most often used metrics to evaluate regression tasks. Another advantage of these metrics is that they can be easily interpreted by domain experts operating the sewer system. We opted for these established

metrics for comparability, reproducibility, and practical applicability.

## 5. Results

Figure 2 shows the distribution of test MSE and MAPE across the 100 random weight initializations for each model type and approach. It is evident in both figures that the LSTM and DeepAR models have a larger spread in the results compared to the other models.



**Figure 2:** Global models often outperform the local models. The distribution of test MSE (upper) and MAPE (lower) across model types and approaches, based on multiple training runs with different random parameter initializations, show that the lowest median MSE values were obtained with LSTM and TFT models. The lowest median MAPE values were obtained with even simpler models such as TCN.


In Table 1, we list the 0.25, 0.5, and 0.75 quantiles ( $q$ ) of the MSE, along with the average measured wall clock runtime at single inference and the effective size of the trained model in Megabytes (MB). We observe that the inference times for global and local models are similar in most cases. This is expected, as the experiments were conducted on high-capacity hardware. It should be

<sup>6</sup><https://optuna.org/>

**Table 1**

Considering MSE, inference time, and model size, the LSTM and TFT models stand out among global and local models. This comparison includes MSE, mean inference time per sample, and effective model size for both global and local approaches. The best values are highlighted in green.

Models	MSE						Inference [ms]		Size [MB]	
	q=0.25		q=0.5		q=0.75		Global	Local	Global	Local
	Global	Local	Global	Local	Global	Local				
LSTM	0.18	0.63	0.24	0.79	0.30	0.99	0.80	0.80	0.31	0.16
TFT	0.23	0.48	0.25	0.50	0.27	0.54	3.08	1.40	11.25	6.22
DeepAR	0.45	1.28	0.61	1.45	0.89	1.64	2.08	2.34	1.71	1.65
Transformer	0.61	0.62	0.61	0.63	0.62	0.64	0.97	1.08	144.26	184.59
N-HiTS	0.69	0.48	0.70	0.48	0.72	0.49	0.92	0.87	1138.41	27.10
TCN	0.98	1.00	0.99	1.01	1.01	1.03	1.24	1.27	0.07	0.04




**Figure 3:** The global LSTM model demonstrates better forecasting performance than the local TFT model. Forecasts from the global LSTM model (see Figure 3a and 3c) and local TFT model (see Figure 3b and 3d) were evaluated on two samples. The plots show the forecasts for the water level of the overflow basin. The dashed red line indicates the start of the forecast window. Samples before the dashed red line represent the input data for the model, spanning 72 hours, while the model produces a 12-hour forecast, shown in blue.

noted that these times were measured while executing the models in parallel on a single GPU.

Among the global models, the LSTM model has the best performance in terms of median MSE. Despite its larger spread in MSE, the LSTM model benefits from

having the second lowest memory consumption and the fastest inference time. The TFT model shows the second-best performance among both global and local models, exhibiting low spread but having the highest measured inference time and relatively high memory consumption.





**Figure 4:** Global models are able to predict sudden changes much better than local models. 12-Hour Ahead Forecasts of the global LSTM (left) and local TFT (right) of Filling Levels Throughout 2023.

In general, the LSTM and TFT models emerge as the most suitable models among the global and local approaches, respectively. However, we highlight the TCN model for its lowest memory consumption, which is five times less than the second-lowest model, the LSTM model. Another notable model is the N-HiTS, which achieves the lowest median MSE among the local models. Considering the memory consumption of the N-HiTS model, it is evident that it requires high-capacity resources for operation.

Figure 3 shows representative examples of forecasts obtained from the LSTM and TFT models. The forecasts shown are from the experiments that achieved the lowest MSE for each model. The global model is represented by the LSTM, while the TFT represents the local model.

In Figure 4, we present an exemplary forecast for a 12-hour horizon into the future. Although exhibiting considerable variability, we observe that the global LSTM model predicts spikes with a higher degree of precision. However, it shows considerable deviations around the mean values near zero. In contrast, the local TFT model struggles to predict sudden changes after longer periods of stagnancy. Evidently, its forecast precision decreases as the forecast horizon extends.

## 6. Conclusion and Future Work

Our results demonstrate that DL models can accurately predict the complex dynamics of wastewater levels in real-world scenarios. Global models, with full access to all sensor readings under normal operation without network outage, exhibit high forecast precision for wastewater levels in the overflow basin. This enhanced precision can significantly aid sewage treatment facilities in effectively redistributing the load of the CSS.

In contrast, local models perform worse in forecasting precision than global models. The reason could be the heavy concentration of target values around the mean.

Due to sudden changes after longer periods of stability, the local models struggle with longer forecasting periods. However, local models can serve as a fallback in the event of a network interruption where exogenous variables become unavailable. Our results indicate that even when all network connections are lost, and only historical sensor readings of an individual sensor are available, adequate forecasts can still be made. Additionally, due to their lower computational costs, it is worthwhile to explore the potential of deploying the local models on edge devices in the future.

## Acknowledgments

The authors gratefully acknowledge the financial support provided by the Federal Ministry for Economic Affairs and Climate Action of Germany for the RIWVER project (Project number: 01MD22007H, 01MD22007C), the Einstein Center Digital Future in Berlin, and the German Research Foundation (Project number: 528483508 - FIP 12).

## References

- [1] S. Bolan, L. P. Padhye, T. Jasemizad, M. Govarathanan, N. Karmegam, H. Wijesekara, D. Amarasingi, D. Hou, P. Zhou, B. K. Biswal, R. Balasubramanian, H. Wang, K. H. Siddique, J. Rinklebe, M. Kirkham, and N. Bolan, "Impacts of climate change on the fate of contaminants through extreme weather events," *Science of The Total Environment*, vol. 909, p. 168388, 2024. doi: 10.1016/j.scitotenv.2023.168388.
- [2] S. F. Thomas J. Wilbanks, "Climate change and infrastructure, urban systems, and vulnerabilities," *And Vulnerabilities: Technical Report for the US De-*

- partment of Energy in Support of the National Climate Assessment, Springer, 2013. doi: 10.5822/978-1-61091-556-4.
- [3] R. Wang, M. J. Eckelman, and J. B. Zimmerman, "Consequential environmental and economic life cycle assessment of green and gray stormwater infrastructures for combined sewer systems," *Environmental science & technology*, vol. 47, no. 19, pp. 11 189–11 198, 2013. doi: 10.1021/es4026547.
- [4] A. Botturi, E. G. Ozbayram, K. Tondera, N. I. Gilbert, P. Rouault, N. Caradot, O. Gutierrez, S. Daneshgar, N. Frison, Ç. Akyol et al., "Combined sewer overflows: A critical review on best practice and innovative solutions to mitigate impacts on environment and human health," *Critical Reviews in Environmental Science and Technology*, vol. 51, no. 15, pp. 1585–1618, 2021. doi: 10.1080/10643389.2020.1757957.
- [5] S. S. Sonone, S. Jadhav, M. S. Sankhla, and R. Kumar, "Water contamination by heavy metals and their toxic effect on aquaculture and human health through food chain," *Lett. Appl. NanoBioScience*, vol. 10, no. 2, pp. 2148–2166, 2020. doi: 10.33263/LIANBS102.21482166.
- [6] M. Schütze, D. Butler, and B. M. Beck, *Modelling, simulation and control of urban wastewater systems*. Springer Science & Business Media, 2002.
- [7] M. M. Saddiqi, W. Zhao, S. Cotterill, and R. K. Dereli, "Smart management of combined sewer overflows: From an ancient technology to artificial intelligence," *Wiley Interdisciplinary Reviews: Water*, vol. 10, no. 3, p. e1635, 2023. doi: 10.1002/wat2.1635.
- [8] T. Züger and H. Asghari, "AI for the public. How public interest theory shifts the discourse on AI," *AI & SOCIETY*, vol. 38, no. 2, pp. 815–828, 2023. doi: 10.1007/s00146-022-01480-5.
- [9] M. Raissi, A. Yazdani, and G. E. Karniadakis, "Hidden fluid mechanics: A navier-stokes informed deep learning framework for assimilating flow visualization data," *CoRR*, vol. abs/1808.04327, 2018. doi: 10.48550/arXiv.1808.04327.
- [10] T. Chiaburu and F. Biessmann, "Interpretable time series models for wastewater modeling in combined sewer overflows," *Procedia Computer Science*, vol. 237, pp. 155–162, 2024. doi: 10.1016/j.procs.2024.05.091.
- [11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.
- [12] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021. doi: 10.1016/j.ijforecast.2021.03.012.
- [13] H. R. Maier and G. C. Dandy, "Neural networks for the prediction and forecasting of water resources variables: a review of modelling issues and applications," *Environmental Modelling & Software*, vol. 15, no. 1, pp. 101–124, 2000. doi: 10.1016/S1364-8152(99)00007-9.
- [14] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," *International journal of forecasting*, vol. 36, no. 3, pp. 1181–1191, 2020. doi: 10.1016/j.ijforecast.2019.07.001.
- [15] C. Challu, K. G. Olivares, B. N. Oreshkin, F. G. Ramirez, M. M. Canseco, and A. Dubrawski, "N-HiTS: Neural hierarchical interpolation for time series forecasting," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 37, 2023. doi: 10.1609/aaai.v37i6.25854. pp. 6989–6997.
- [16] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan, and L. Sun, "Transformers in time series: A survey," *arXiv preprint arXiv:2202.07125*, 2022. doi: 10.48550/arXiv.2202.07125.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017. doi: 10.48550/arXiv.1706.03762.
- [18] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018. doi: 10.48550/arXiv.1803.01271.
- [19] T. Ling, J. Hoefer, C. Qian, and G. Schiele, "Flowprecision: Advancing fpga-based real-time fluid flow estimation with linear quantization," in *2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. IEEE, 2024. doi: 10.1109/PerComWorkshops59983.2024.10503436. pp. 733–738.
- [20] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. doi: 10.48550/arXiv.1412.6980.
- [21] H. Hewamalage, P. Montero-Manso, C. Bergmeir, and R. J. Hyndman, "A look at the evaluation setup of the m5 forecasting competition," *arXiv preprint arXiv:2108.03588*, 2021. doi: 10.48550/arXiv.2108.03588.

## A. Hyperparameter Optimization

Hyperparameter optimization was performed using the Tree-structured Parzen Estimator algorithm provided in the *Optuna* library<sup>7</sup>. The optimization process was conducted in two iterations with a total of 600 trials:

1. **Broad Search with 500 trials:** An extensive hyperparameter search space will be explored to identify potential optimal values.
2. **Refined Search with 100 trials:** Based on the results of the broad search, a more focused and fine-grained search will be conducted around the best-performing hyperparameters.

The hyperparameters optimized on the validation data include several hyperparameters shared by all models and some model-specific hyperparameters.

After the first iteration of fine-tuning, the input and prediction sequence lengths were set to 72 hours and 12 hours, respectively. Additionally, the batch size was set to 256, as this configuration worked well across all models. Below, we list the optimal hyperparameter values obtained after the second iteration for both the local and global models:

**Table 2**  
Optimal Hyperparameter Settings for LSTM Models

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	0.0084	0.0228
batch_size	256	256
weight_decay	4.1575e-05	1.7515e-04
dropout	0.2169	0.2895
hidden_dim	36	33
n_rnn_layers	1	1

**Table 3**  
Optimal Hyperparameter Settings for DeepAR model

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	0.0460	0.0295
batch_size	256	256
weight_decay	2.004e-05	1.110e-05
dropout	0.3142	0.3776
hidden_dim	128	106
n_rnn_layers	1	1

**Table 4**  
Optimal Hyperparameter Settings for N-HiTS model

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	8.506e-05	2.611e-04
batch_size	256	256
weight_decay	4.014e-04	6.409e-03
dropout	0.3027	0.5367
num_stacks	3	2
num_blocks	5	5
num_layers	1	1
layer_widths	1024	1024

**Table 5**  
Optimal Hyperparameter Settings for Transformer model

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	4.642e-05	2.366e-05
batch_size	256	256
weight_decay	1.489e-02	1.141e-02
dropout	0.07572	0.05084
num_encoder_layers	3	2
num_decoder_layers	6	3
d_model	96	132
nhead	3	3
dim_feedforward	4096	4096

**Table 6**  
Optimal Hyperparameter Settings for TCN model

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	0.0436	0.0359
batch_size	256	256
weight_decay	0.0162	0.0268
dropout	0.4212	0.3418
num_filters	4	4
dilation_base	5	5
kernel_size	4	3
weight_norm	False	True

**Table 7**  
Optimal Hyperparameter Settings for TFT model

Hyperparameters	Optimal Values	
	Local Model	Global Model
learning_rate	0.0025	0.0041
batch_size	256	256
weight_decay	1.575e-05	1.845e-05
dropout	0.2666	0.1087
hidden_continuous_size	14	17
hidden_size	54	54
lstm_layers	3	2
num_attention_heads	2	3
full_attention	True	False

<sup>7</sup><https://optuna.org/>