


ZTF2026

HPC-ENABLED BILSTM ARCHITECTURE FOR MULTI-STEP
ENERGY CONSUMPTION FORECASTING

Jelena Joksimović, Martin Zagorc
Rudolfovo Science and Technology Center

Science and Technology Forum 2026, 31.3.2026



Introduction

- Buildings amongst the largest energy consumers
- Migration from static control to adoptive, data-driven energy management
- To be able to accurately predict the electricity consumption for **the next day** with **15-minute granularity**
- Data: from the meter -> past consumption with 15-minute granularity + weather data

Introduction

- Collaboration between Rudolfovo (Program for Energy Management) and Robotina
- Fortissimo Plus Project - EuroHPC: Advanced Multi-objective optimization of Energy management systems using Deep neural networks and HPC for real-time, multi-step energy forecasting (AIMED-HPC)


State-of-the-art

- Liu, C., Zhang, C.: A day-ahead prediction method for high-resolution electricity consumption in residential units. *Energy*. 265. (2023)
- Li, Z., Yu, F., Guan, Q.: Bilstm network-based approach for electric load forecasting in energy cell-tissue systems. *2022 First International Conference on Cyber-Energy Systems and Intelligent Energy (ICCSIE)*. pp. 1–5 (2023).
- Binbusayyis, A., Sha, M.: Energy consumption prediction using CNN-BiLSTM. *Heliyon* 10 (2024)
- Lian H. et al.: Research on Real-Time Energy Consumption Prediction Method and Characteristics of Office Buildings Integrating Occupancy and Meteorological Data. *Buildings* 2025, 15, 404. (2025)

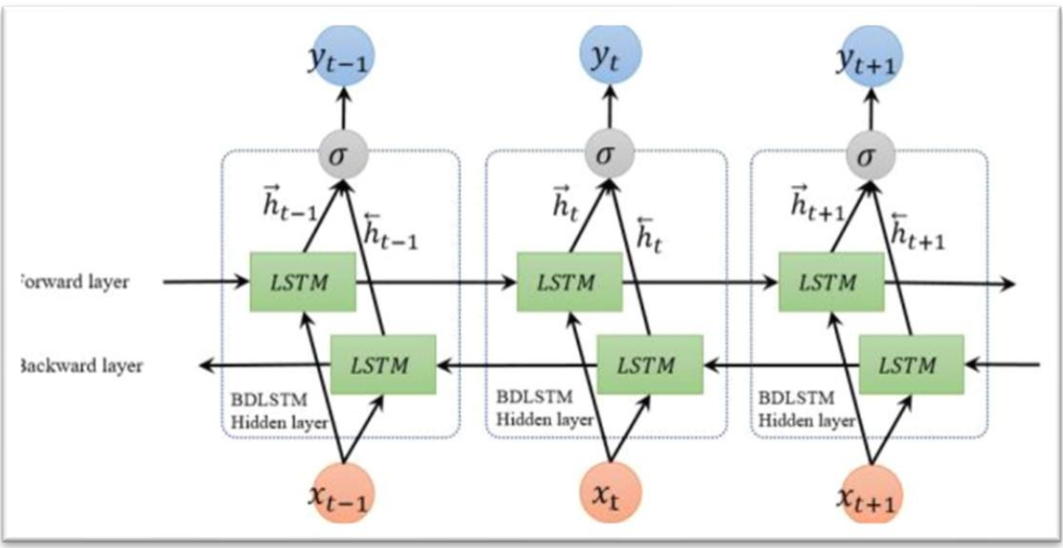
BiLSTM Neural Networks

BiLSTMs process information in both forward and backward directions, offering a comprehensive view of the input data's temporal dynamics

[1] The Performance of LSTM and BiLSTM in Forecasting Time Series. S. Siami-Namini, N. Tavakoli & A. Namin (2019). *2019 IEEE International Conference on Big Data*.



LSTM memory cell



Structure of BiLSTM neural network

BiLSTM Neural Networks bottlenecks

- Computationally expensive
- Standard hardware is unsuitable for large-scale deployment
 - in our case, that means to predict consumption of electricity for 1 day in advance in a 15-minute granularity for 1,000+ buildings
- We need to train a model on time series data with long sequences as input and a multistep vector as an output (96)

Neural Architecture Search

- In this article, as a *Task 3* of the *FFPlus project*, we run a Neural Architecture Search (NAS) on HPC, because this task was infeasible on Robotina's or our local hardware.
- Example is 298-long input vector and 96-long output vector, with 30 features (time components, weather, lags of energy consumption), that leads to **11h+ per trial** on the local machine for only one building
- Leading to a **900+ hours** to do a full NAS testing hundreds of parameters

BiLSTM Neural Networks

- This is exactly what our AIMED-HPC project aims for:
 1. Leverage HPC infrastructure to search for and find the best parameters for the training of BiLSTM models for accurate electricity consumption prediction for buildings
 2. Leverage Teacher-Student networks to do a knowledge distillation - Large BiLSTM models trained on HPC infrastructure will act as teacher models, while smaller, compressed student models will be distilled to maintain accuracy with significantly reduced complexity. This will enable efficient deployment of predictive models on resource-constrained devices.

VEGA HPC Cluster

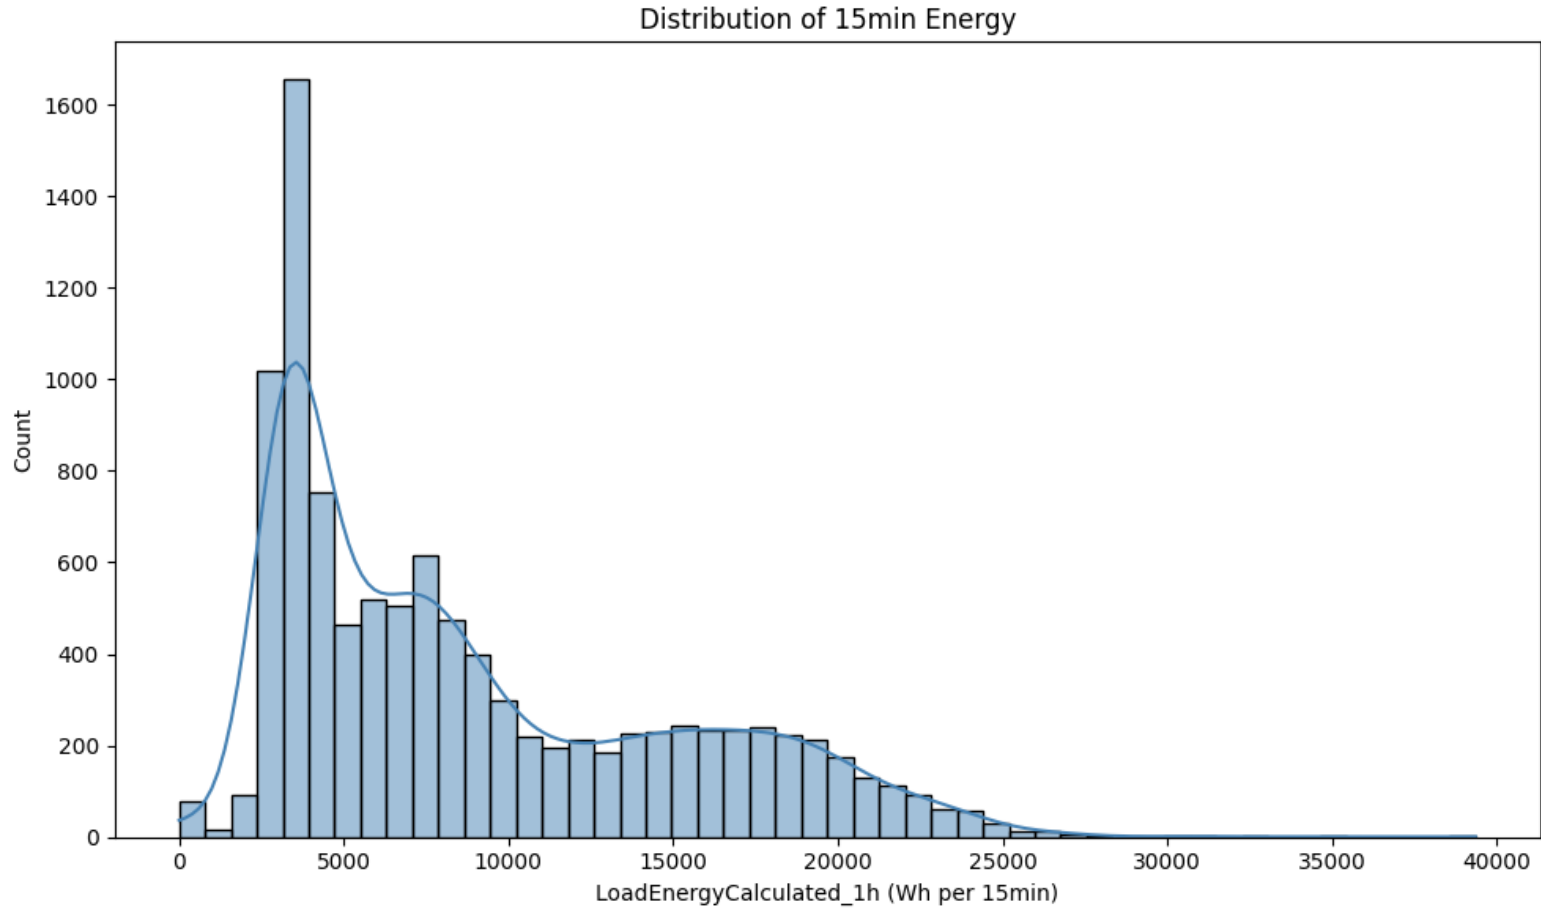
HPC Resources

- 1,000 node-hours and 128,000 GPU core-hours annually
- 3 GPUs per training job (TensorFlow 2.15.0 in Singularity containers)

End-to-end workflow

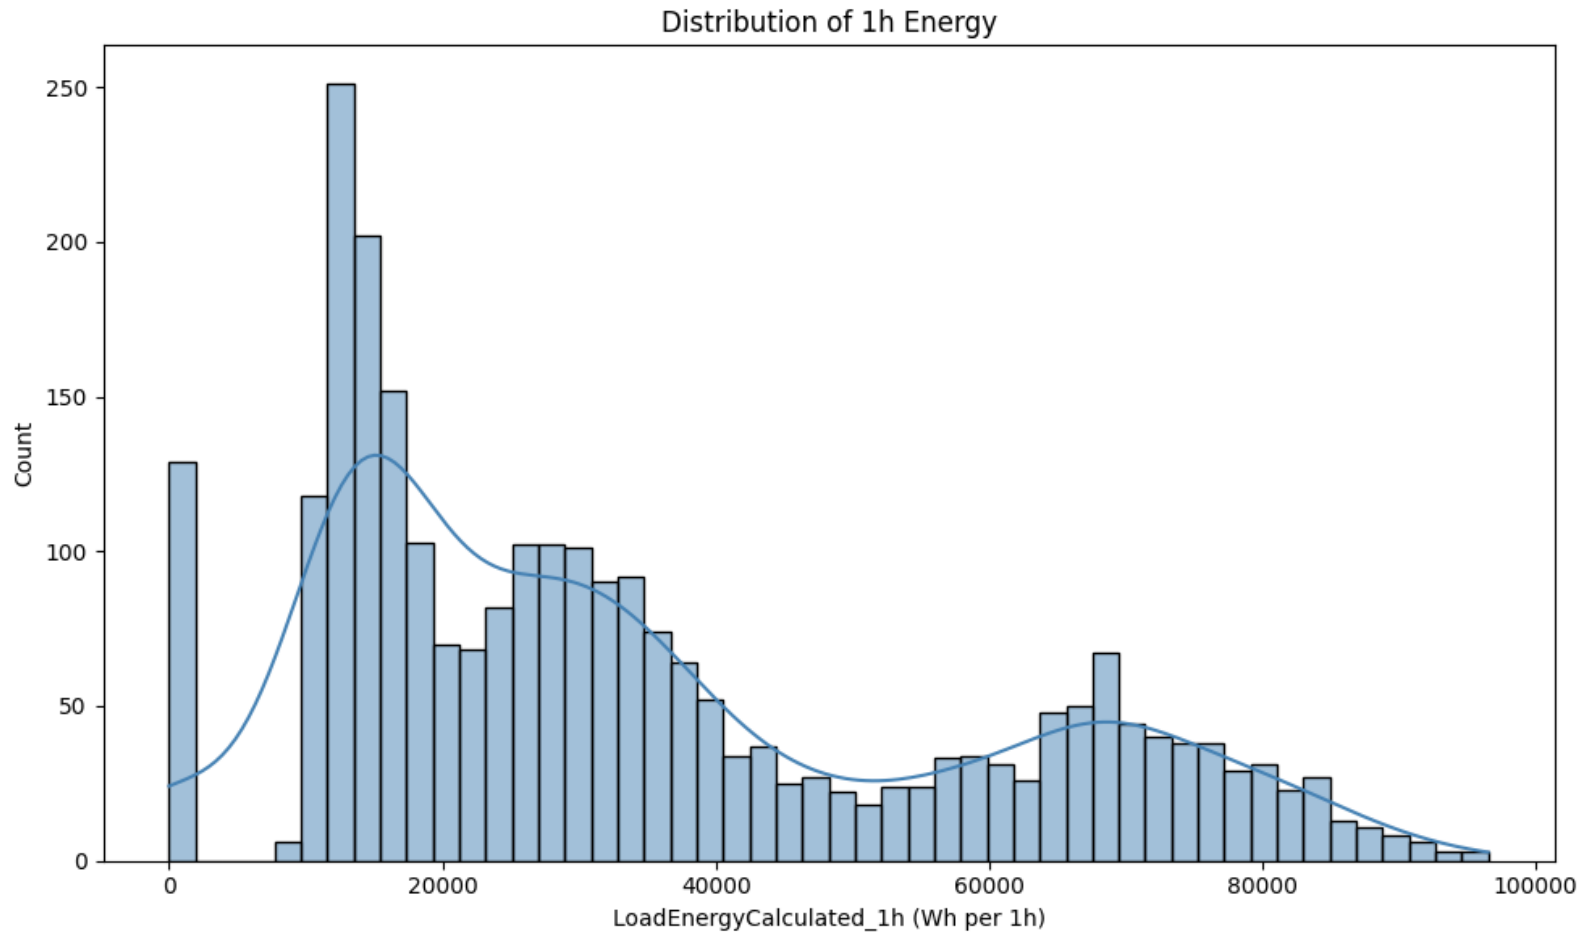
- Automated Data Ingestion: Smart meter data (MQTT), weather (OpenWeatherMap / Open-Meteo), market prices, tariff rules in Slovenia
- Structured Storage: PostgreSQL + TimescaleDB for efficient time-series operations
- Remote HPC Training: SLURM batch jobs via SSH with secure API access
- Efficient Synchronization: Monthly retraining cycles with consistency checks
- Deployment Automation: Models & scalars retrieved via PostgreSQL triggers, deployed into Robotina's EMS for real-time optimization

Data – electricity consumption per 15-min



count	10231.000000
mean	8914.137885
std	6019.352009
min	0.000000
25%	3713.704545
50%	7103.454545
75%	13287.666666
max	39360.888888

Data – electricity consumption per 15-min



count	2672.000000
mean	34151.818959
std	23340.193873
min	0.000000
25%	14962.333334
50%	27887.111111
75%	50270.352273
max	96523.000000

Feature Selection

- season (from 0 to 3)
- energy consumption lags (from 1 to 21)
- month (from 1 to 12)
- day of the week (from 1 to 7)
- is weekend? (0 or 1)
- is holiday? (0 or 1)
- temperature
- humidity
- temperature * hour
- hour (cos transformation)

Results

Table 1: Parameters explored during NAS

Parameter type	Values tested
Input context: sequence length (timesteps)	298, 384, 480, 596, 672
Model hyperparameter: units per BiLSTM layer	64, 128 , 256
Model hyperparameter: number of BiLSTM layers	1, 2
Model hyperparameter: dropout rate	0.0 , 0.2, 0.4
Model hyperparameter: learning rate	1×10^{-3} , 5×10^{-4} , 1×10^{-4}

Each NAS trial dynamically constructed a BiLSTM model by exploring different combinations of key hyperparameters. The search was carried out using Keras Tuner's Hyperband algorithm, which systematically prunes underperforming models and allocates more resources to promising ones. The best-performing configuration (highlighted in yellow) was selected based on the lowest validation Mean Absolute Error (MAE).

Results

Table 5: Performance comparison between baseline and BiLSTM model

Metric	Baseline model	BiLSTM model
MAE	256.27	170.79
RMSE	408.40	608.40
SMAPE (%)	47.88	10.49
R^2	0.39	0.73

Future work

- Train large BiLSTM models on HPC as teachers; distill into lightweight student models for deployment on smart meters
- Enable real-time forecasting across thousands of buildings with reduced model complexity on resource-constrained devices.
- Investigate Transformer-based architectures to better handle long temporal dependencies and leverage parallelization.

ZTF2026

Thank you for your attention!

Questions?

Jelena Joksimović

jelena.joksimovic@rudolfovo.eu

Rudolfovo - Science and Technology Centre Novo mesto


RU
DOL
FO
VO

