

Continual Learning for very Short-Term Load Forecasting: A Case Study on Parts Cleaning

Project Manager
Robin Zink

Principal Investigator
Prof. Dr.-Ing. Matthias Weigold

Project Term
2024 - 2025

Clusters
Lichtenberg II Cluster Darmstadt

Software
TensorFlow, PyTorch

Institute
Production Management, Technology
and Machine Tools

University
Technische Universität Darmstadt



Introduction

Industrial production systems play a central role in transition toward more flexible and sustainable energy use. In particular, energy-intensive machines must increasingly adapt their electricity consumption to volatile power supply conditions. Reliable very short-term forecasts of electrical load are therefore essential to enable informed operational decisions. However, industrial environments are highly dynamic. Machines are reconfigured, operating conditions change and external influences vary over time. These changes cause data patterns to evolve continuously, which makes static prediction models unreliable once deployed. The project addressed this challenge by developing adaptive prediction models that can learn continuously from new data while maintaining previously acquired knowledge. The project evaluated multiple continual learning approaches on a time series dataset with various manually induced data drifts. HPC resources were essential to handle the computational demands arising from repeated model training over multiple seeds, with various hyperparameter configurations. Without HPC, the experimental scale achieved would not have been feasible.

Methods

The project implemented a two-stage workflow. First, an initial prediction model was trained using historical data. Historical sensor data describing machine operation and electricity

consumption were processed and used to train neural network models capable of capturing temporal dependencies. Second, during simulated deployment, the model performance was continuously monitored. When significant deviations in prediction quality occurred, indicating changes in the underlying process, the model was updated using continual learning approaches that balance out adaptation to new data with preservation of earlier knowledge. This work focused exclusively on regularization-based methods, which aim to balance the preservation of previous knowledge and learning new knowledge by regularizing the loss term. Several continual learning approaches were implemented and compared under identical conditions. Their performance was assessed against two baseline approaches. The first baseline relies on a static model that is trained once and then deployed without any further updates. The second baseline includes performance monitoring and retraining, but updates the model using new data only, without any mechanisms to preserve previously learned knowledge.

Results

The project produced a published case study [1] and a framework for continual load forecasting in industrial environments. The HPC resources enabled a study of 260 experiments. The experimental results demonstrate that continual learning approaches consistently outperformed both baseline methods in terms of predictive accuracy over time. In particular, the continual learning models were quick to stabilize their performance after changes in operating conditions, while the static baseline showed a pronounced degradation. Compared to simple retraining without knowledge preservation, the continual learning approaches achieved markedly better stability and retained previously learned behavior, leading to more reliable forecasts throughout the entire deployment phase.

Discussion

The results confirm that continual learning strategies are well suited for short-term load forecasting in dynamic industrial environments. Static models showed a clear decline in prediction quality once operating conditions changed, demonstrating their limited suitability for long-term deployment. Simple retraining improved adaptability but often led to unstable performance due to the loss of previously learned behavior. In contrast, the continual learning approaches maintained a better balance between adaptation and stability. They reacted quickly to changes in operating conditions while preserving knowledge from earlier operating phases, resulting in more reliable predictions over time. This stability is particularly important in industrial applications, where inconsistent model behavior can limit practical usability. The extensive experimental evaluation relied heavily on HPC resources. The large number of experiments across different learning strategies and parameter settings was essential to ensure robust and statistically meaningful results.

Outlook

Future work includes extending the framework to additional continual learning strategies, such as rehearsal-based methods, which were not covered in the present study. Furthermore, additional use cases are planned for validation

Publications

[1] Zink, R.; Magin, J.; Griess, O.; Weigold, M.: "Continual learning for very short-term load forecasting: A case study on parts cleaning", Applied Energy, Volume 402, Part A, 2025, 126905
<https://doi.org/10.1016/j.apenergy.2025.126905>

Last Update: 2026-03-31 10:16