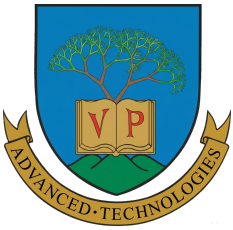
University of Pannonia
Faculty of Information Technology
Doctoral School of Information Science and Technology




# Prediction, Control, and Design: Machine Learning and Optimisation Methods for Battery-Enhanced Renewable Energy Systems

Thesis booklet

Author: Tamás Miseta

Supervisors:
Ágnes Vathy-Fogarassy PhD
Attila Fodor PhD

Veszprém
2026

# 1 Motivation and research objectives

The global transition toward weather-dependent renewable energy sources (RES) has introduced complex, interconnected challenges in electricity market forecasting, the predictive control of battery-enhanced power plants, and the physical design of photovoltaic (PV) systems. In competitive electricity day-ahead markets (DAM), producers and consumers submit bids for the delivery blocks of the following day (trading strategy); the clearing price is determined by the intersection of aggregated supply and demand curves [1]. The increasing share of non-dispatchable wind and solar generation has made these prices more volatile, creating opportunities for battery-enhanced PV power plants to shift energy between cheap and expensive hours, provided prices can be predicted accurately [1]. With a suitable power-shifting strategy, the economic return on renewable projects can be increased, thus contributing to the wider adoption of weather-dependent renewables in the energy mix [2]. These strategies can also contribute to grid stability and reduce the $CO_2$ impact of the energy mix.

To achieve these goals, three classes of methods are needed: prediction methods that provide reliable, accurate insights into day-ahead electricity market dynamics. Control methods for RES-based microgrids that rely on predicted electricity prices and a system model to develop trading strategies. Sizing methods that evaluate and optimise physical system configurations based on the results of the trading strategies.

Although existing methods have reached a certain level of maturity, important unresolved challenges still motivate the research presented in this dissertation:

(i) Early stopping, a widely used stopping criterion in machine learning [3], monitors only the validation loss during training, which may be suboptimal. By jointly analysing the dynamics of training and validation loss, more reliable detection of overfitting becomes possible.

(ii) Long Short-Term Memory (LSTM) models [4] are widely used in time-series prediction, with great success [5, 6]. However, they could still be improved with new regularisation techniques, large-scale hyperparameter optimisation, and cross-validation.

(iii) A new class of foundation time-series prediction models emerged [7, 8, 9]: Transformer-based models that were pre-trained on a large corpus of time-series data. Although there is limited preliminary research on these models [10], the context-length sensitivity of time-series foundation models for electricity price forecasting (EPF) remains largely unexplored.

(iv) Although there are countless battery energy storage system models present in the literature [11], there are few that are comprehensive enough to accurately

model the behaviour of these systems and efficient enough to be used in large-scale simulations involving system size and trading strategies. Therefore, the development of a simplified system model with the aforementioned qualities became a possibility.
(v) The control of microgrids that contain a battery energy storage system, a photovoltaic system, and a consumer are studied extensively in the literature [11]; however, they could be improved from economic, sustainability, and engineering aspects. These strategies can contribute to grid stability, a topic that could be studied in greater depth.
(vi) Currently, the photovoltaic design methods in the literature are rather empirical than systematic [12], and there are few automatic multi-string PV energy system design algorithms. The development of an algorithm that optimises multi-string PV energy systems under feed-in restrictions, based on consumption data and with per-string convergence guarantees for verifiably unimodal objectives, became possible.
(vii) Current methodologies for the physical design of microgrids containing a semi-autonomous residential consumer, a multi-string PV energy system could be improved. The development of an energy system design methodology for semi-autonomous structures involving a multi-string PV energy system and a battery became possible.

My research addresses these gaps along three interconnected pillars, formulated as three research objectives:

**Objective 1.** *To develop machine learning methods for accurate day-ahead electricity price forecasting (EPF).*

**Objective 2.** *To develop optimisation algorithms and system models for energy trading strategies that maximise the profit of battery-enhanced renewable power plants using predicted electricity prices.*

**Objective 3.** *To develop optimisation algorithms and system modelling methodologies for the physical design of PV energy systems under real-world operational, regulatory, and sustainability constraints.*

# 2 Scientific results formulated as theses

## 2.1 Machine learning prediction methods for day-ahead electricity price forecasting

LSTM networks [4] have been widely and successfully applied to electricity price forecasting [5, 13] and load forecasting [6]. A key challenge in training such models is overfitting: training is typically halted by early stopping when the validation loss fails to improve for a specified number of epochs [3]. However, this approach monitors only a single loss curve and cannot distinguish a genuine plateau from a temporary fluctuation. As a result, it may lead to premature stopping or unnecessarily prolonged training. To address this limitation, a novel stopping criterion was developed that jointly analyses training and validation loss dynamics during the training of neural network models.

Recently, pre-trained Transformer-based time-series foundation models, including Chronos [7, 14], TimesFM [8], and Moirai [9], have emerged as a new class of models that leverage self-attention [15] and large-scale pre-training. Early benchmarks [10] indicate that foundation models perform on par with traditional EPF approaches but do not consistently outperform well-tuned task-specific models; furthermore, their context-length sensitivity in EPF remained unexplored. These observations motivate a systematic investigation of foundation models and their behaviour under varying context lengths in EPF tasks.

### Thesis I.1: Correlation-Driven Stopping Criterion for neural network training

I developed the Correlation-Driven Stopping Criterion (CDSC) for stopping neural-network training at the onset of overfitting, detected from the rolling Pearson correlation between the training and validation loss. Across four benchmark problems and three network architectures, it matched or exceeded both early stopping and the maximum-number-of-epochs criterion, with clear gains even on problems where those conventional criteria offered little or no improvement. The findings were supported by a comparative regret-bound analysis and a parameter-sensitivity study.

*Publication: [P1]*

### Thesis I.2: LSTM-based day-ahead electricity price forecasting with CDSC-enhanced training

I developed hyperparameter-optimised univariate Long Short-Term Memory (LSTM) models for day-ahead electricity price forecasting, first for the Hungarian (HUPX) market and subsequently in a cross-market setting covering the Germany, Singapore, and United Kingdom day-ahead markets, with the CDSC tuned jointly with the network. The optimisation covered the input lookback, the network architecture, and the regularisation settings. In the cross-market experiments, the CDSC-trained models matched or outperformed conventionally early-stopped models in 10 of the 12 investigated configurations (three markets, two prediction horizons, and two validation schemes), with the largest gain, a 16.7% test mean absolute error (MAE) reduction, at the longer horizon on the volatile Singapore NEMS market.

*Publications: [P2], [P5]*

### Thesis I.3: Transformer-based foundation models for direct multi-step electricity price forecasting

I developed a systematic, context-length-optimised zero-shot benchmark of six pre-trained Transformer foundation models for direct multi-step day-ahead electricity price forecasting, evaluated against the hyperparameter-optimised univariate LSTM models and conventional forecasting baselines. In all six market-horizon configurations the best foundation model attained a lower test MAE than the optimised LSTM despite using no task-specific training, reducing it by 18.4–40.8%, and Chronos 2 ranked first in five of the six. The ranking held under Diebold–Mariano tests, robustness tests, and an independent holdout year, though the foundation models benefit from large-scale pre-training that the univariate LSTM cannot access.

*Publication: [P5]*

## 2.2 Energy trading strategy optimisation for battery-enhanced renewable power plants

The economic viability of most battery-enhanced renewable power plants depends on the trading strategy employed on the day-ahead market [2]. Differential evolution (DE) [16] has been applied in electric power systems for economic dispatch [17] and renewable energy system sizing [18], providing a flexible, gradient-free optimisation approach. However, domain-specific methods that explicitly exploit the structure of the energy trading problem, as well as parametrised system models capable of capturing the core physical constraints of DC-coupled systems in a continuous formulation, remain limited.

These observations motivate the development of both structure-exploiting gradient-based optimisation methods and robust gradient-free alternatives for the optimisation of day-ahead electricity trading strategies based on predicted electricity prices.

**Thesis II.1: Adaptive gradient-descent-based energy trading optimisation**

I developed Optimised Energy Trading by Gradient Descent (OptEnT-GD), a novel adaptive gradient-descent method for optimising the 24-hour day-ahead market (DAM) trading strategies of battery-enhanced photovoltaic (PV) power plants from predicted electricity prices, under battery-capacity and inverter-power constraints. I evaluated the method across 42 system configurations covering three inverter powers, seven battery capacities, and two storage efficiencies. The daily profit gain over immediate dispatch grew with the battery-capacity-to-inverter-power ratio and, on the configurations with 5% storage loss, reached 52.94% with perfect price information and 46.26% with LSTM-based price forecasts; the prediction error reduced the profit increase by only 1.93–6.93 percentage points relative to the perfect-information case.

*Publication: [P2]*

**Thesis II.2: Differential-evolution-based energy trading optimisation and system sizing guide**

I developed Optimised Energy Trading by Differential Evolution (OptEnT-DE), a derivative-free method for 24-hour DAM trading-strategy optimisation that imposes no initial-monotonicity assumption. It pairs the standard differential-evolution operators with a problem-specific formulation, in which each candidate solution encodes a complete 24-hour trading schedule and the fitness is the constrained profit objective, and it therefore accommodates, by construction, a broader class of nonlinear, non-differentiable objectives. Benchmarked against OptEnT-GD on the same 42 configurations under predicted prices, the best of 100 daily runs reached a maximum profit approximately 3% below that of OptEnT-GD at the largest battery capacity, the gap growing with battery size. The analysis across the configurations further showed that the relative profit gain increases with the battery-capacity-to-inverter-power ratio, providing an empirical basis, as a system sizing guide, for selecting the storage capacity in relation to the inverter power.

*Publication: [P2]*

## 2.3 Optimisation and modelling of physical energy systems for industrial and household-size prosumers

Industrial and household-size prosumers face diverse constraints that render conventional PV design methods suboptimal [12]. The classical bisection (dichotomous search) method [19] guarantees convergence for unimodal objectives but has not been adapted to consumption-aware multi-string PV optimisation under feed-in restrictions. For semi-autonomous structures (operating independently of the public grid for most of the year while retaining a limited grid backup), an integrated methodology was needed to combine load estimation, thermal modelling, and PV-storage co-optimisation. The majority of the methods used in industry are more empirical than systematic. There are two methods widely used in the industry: the first, the *Sum Yearly Consumption Matching* (SYCM) method, sizes the PV energy system so that its estimated yearly output equals the yearly energy consumption of the site. The second method limits the energy lost to feedback protection to a predefined threshold ($E_{\text{lost}}$ *limit method*). These methods are not automatic and do not produce optimal systems from the economic, engineering, or sustainability standpoint. The contributions of this thesis group lay the foundations for the systematic and automatic design of systems that are optimal from engineering, economic, and sustainability standpoints.

### Thesis III.1: Consumption-aware dichotomous search-based string optimisation of multi-string photovoltaic systems

I developed the dichotomous search-based string optimisation (DSSO) algorithm, a consumption-aware method for designing multi-string PV energy systems of industrial prosumers who cannot feed surplus energy into the grid. For each string it returns the optimum of any objective unimodal in the output power of that string, to a single-panel tolerance in a logarithmic number of steps; multi-string systems are treated by coordinate-wise passes over the strings in order of their sensitivity to the objective, converging to a fixed point that is optimal against single-string perturbations. Under a fixed tariff, the optimised systems achieved 14.8–17.5% lower grid consumption, 2.7–3.2 times higher self-reliance, and 12.9–15% lower $CO_2$ emissions than the $E_{\text{lost}}$ limit method, at only 37–45% of the system size returned by the Sum Yearly Consumption Matching method; under a DAM-based tariff the method remained applicable and returned smaller optimal systems.

*Publication: [P4]*

### Thesis III.2: Design methodology for semi-autonomous electric energy systems

I developed a design methodology for the energy systems of semi-autonomous floating homes, combining the Stochastic Consumption Estimation Algorithm (SCEA), a stochastic load generator that requires no site-specific metered consumption data, with a coupled thermal-electrical system model and discrete enumeration over PV and battery energy storage system (BESS) configurations under zero feed-in and a weight-limited battery. In a year-long case study of a real floating-home design covering 12 PV–BESS configurations, the largest PV array (26.65 kWp across three orientations) with a 20 kWh battery reached 91.43% self-reliance and an 86.2% reduction in grid consumption, but left 22.14 MWh of generation unutilised per year. A roof-only array with the same battery retained 89.59% self-reliance and an 83.7% reduction while lowering the unutilised energy to 16.53 MWh, so the smaller, roof-only system proved the more effective design choice.

*Publication: [P3]*

# 3 Application of the results and future directions

The methods developed in this dissertation have practical applications in three interconnected domains, and each leaves concrete directions for further work.

The CDSC is applicable to any supervised neural network training process in which the validation loss correlates nontrivially with the training loss, and extends beyond the energy domain; its architecture-agnostic formulation makes it a drop-in replacement for early stopping. Its three parameters (window, threshold, and patience) span a larger search space than that of early stopping, but Bayesian optimisation with gradient-boosted regression trees located suitable settings within 35 to 60 evaluations; the best settings are task-dependent and are therefore tuned per problem. Its integration with the PyTorch and TensorFlow frameworks as a callback, and its application to further forecasting tasks such as PV generation and load forecasting, are planned.

The LSTM-based prediction pipeline and the foundation model evaluation provide deployable forecasting solutions for day-ahead electricity markets at 30- and 60-minute resolution, given several years of historical prices. The finding that zero-shot foundation models can substantially outperform trained LSTMs is particularly relevant for market operators requiring accurate forecasts without the overhead of task-specific training, provided that the available history spans the longest tested context length. The model rankings behind Theses I.2 and I.3 were tested for statistical significance with the Diebold–Mariano test [20] and confirmed with Model Confidence Sets on a held-out year. The demonstrated cross-market applicability, spanning the Hungarian day-ahead market (HUPX), the Germany day-ahead market (EPEX SPOT), the Singapore day-ahead market (NEMS), and the United Kingdom day-ahead market (N2EX), suggests generalisation across structurally distinct markets, although the optimal CDSC threshold proved market-specific (0.0 to 0.7) and the foundation models were evaluated zero-shot only, without fine-tuning on DAM data. Future work will extend the study to further markets and intermediate horizons, fine-tune the leading foundation models on DAM corpora, and add quantile or distributional outputs that propagate forecast uncertainty into risk-aware bidding.

The OptEnT-GD and OptEnT-DE algorithms, combined with the parametrised DC-coupled system model, provide a complete workflow for developing trading strategies for battery-enhanced renewable power plants on hourly-resolution DAMs with 24-hour granularity and standard gate closure. It can be applied to household-sized systems, virtual power plants (VPPs), and energy communities. OptEnT-GD applies whenever the trading objective is continuous, differentiable, and initially monotone in the step size; this held for the profit objective studied, but it may fail once battery degradation, transaction costs, detailed inverter

losses, or non-convex tariffs are added, in which case OptEnT-DE serves as the derivative-free fallback, at the price of a 25 to 30% spread between the best and worst of 100 runs, so that a single run gives no profit guarantee. The system sizing guide provides practitioners with an empirical basis for selecting the storage capacity in relation to the inverter power at the procurement stage. Both methods rest on a simplified DC-coupled model, so the main direction for future work is to validate the model on a real-world system and then to extend the objective with more complex tariff structures and with economic, engineering, and sustainability metrics beyond profit, while reducing the stochastic spread of OptEnT-DE.

The DSSO algorithm and consumption-aware evaluation framework provide an optimisation tool for industrial prosumers operating under feed-in restrictions wherever reliable load data are available at sub-hourly resolution, converging in logarithmic time with any objective unimodal in the string power and any tariff structure that preserves that unimodality; a new metric requires the same one-dimensional unimodality check before use. The benchmark is scoped to the standard engineering sizing rules (SYCM and the $E_{\text{lost}}$ limit method); a comparison with optimisation-based sizing tools such as HOMER, PVSyst workflows, and mixed-integer linear programming formulations; metaheuristic handling of non-unimodal objectives; a broader set of economic, engineering, and sustainability metrics; and validation on a wider corpus of industrial load profiles are planned. The comprehensive methodology for semi-autonomous structures (SCEA, thermal-electrical model, joint PV-storage optimisation) enables designers to evaluate configurations before construction, avoiding costly over- or under-sizing, wherever a building targets minimal grid reliance under tight weight, orientation, and thermal-mass constraints and appliance-level usage data substitute for metered consumption; the thermal coupling is calibrated for water-surrounded structures but may generalise to lightweight prefabricated dwellings. As the validation is simulation-based and covers a single case study, future work will focus on the physical validation of the completed floating home, on the end-to-end coupling of the thermal and electrical subsystems within the optimisation loop, and on anchoring the SCEA profiles against metered consumption from comparable structures.

Across the three groups, the methods were matched to the structure of each subproblem, which offers a reusable selection guideline: the CDSC regularises training wherever the training and validation losses correlate informatively, OptEnT-GD exploits the initial monotonicity of the profit objective, and DSSO relies on the unimodality of the objective in the string power; where such structure is absent, population-based (OptEnT-DE) or pre-trained general-purpose (foundation) models take over.

The long-term potential is a unified decision-support workflow integrating price forecasting (Thesis I), energy trading optimisation (Thesis II), and physical system design (Thesis III) into a single tool for prosumers and energy communities. Its realisation, through the coupled validation of the three stages on a single day-ahead market and site so that forecasting, trading, and design reinforce one another, is the principal direction for future work.

# 4 Publications related to the dissertation

**P1** Tamás Miseta, Attila Fodor, and Ágnes Vathy-Fogarassy. Surpassing early stopping: A novel correlation-based stopping criterion for neural networks. *Neurocomputing*, 567:127028, 2024. **[Q1, IF: 6.7]**

**P2** Tamás Miseta, Attila Fodor, and Ágnes Vathy-Fogarassy. Energy trading strategy for storage-based renewable power plants. *Energy*, 250:123788, 2022. **[D1, IF: 10.1]**

**P3** Tamás Miseta, Ágnes Vathy-Fogarassy, Zsófia Závodi-Fodor, and Attila Fodor. Design, modeling, and simulation of the energy system of semi-autonomous floating homes. In *29th International Conference on Intelligent Engineering Systems (INES 2025), Palermo, Italy*, pages 143–148. IEEE, 2025.

**P4** Tamás Miseta, Ágnes Vathy-Fogarassy, Zsófia Závodi-Fodor, and Attila Fodor. Consumption-aware optimization of multi-string photovoltaic systems based on economic metrics. In *5th International Conference on Electrical, Computer and Energy Technologies (ICECET 2025), Paris, France*, pages 1–6. IEEE, 2025.

**P5** Tamás Miseta, Attila Fodor, and Ágnes Vathy-Fogarassy. Electricity price prediction with univariate Long Short-Term Memory and zero-shot foundation transformer models. *Energy and AI*, 100886, 2026. **[D1, IF: 9.7]**

# References


[1] Rafał Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4):1030–1081, 2014.

[2] Muhammad Khalid, Ricardo P Aguilera, Andrey V Savkin, and Vassilios G Agelidis. On maximizing profit of wind-battery supported power station based on wind power and energy price forecasting. *Applied Energy*, 211:764–773, 2018.

[3] Lutz Prechelt. Early stopping – but when? In Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade*, volume 7700 of *Lecture Notes in Computer Science*, pages 53–67. Springer, 2 edition, 2012.

[4] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[5] Umut Uğurlu, Ilkay Oksuz, and Oktay Tas. Electricity price forecasting using recurrent neural networks. *Energies*, 11(5):1255, 2018.

[6] Salah Bouktif, Ali Fiaz, Ali Ouni, and Mohamed Adel Serhani. Optimal deep learning LSTM model for electric load forecasting using feature selection and genetic algorithm: comparison with machine learning approaches. *Energies*, 11(7):1636, 2018.

[7] Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.

[8] Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. *arXiv preprint arXiv:2310.10688*, 2024. TimesFM 2.0.

[9] Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Unified training of universal time series forecasting transformers. *arXiv preprint arXiv:2402.02592*, 2024.

[10] Timothée Hornek, Amir Sartipi, Igor Tchappi, and Gilbert Fridgen. Benchmarking pre-trained time series models for electricity price forecasting. In *Proceedings of the 2025 21st International Conference on the European Energy Market (EEM)*. IEEE, 2025.

[11] Shaopeng Guo, Qibin Liu, Jie Sun, and Hongguang Jin. A review on the utilization of hybrid renewable energy. *Renewable and Sustainable Energy Reviews*, 91:1121–1147, 2018.

[12] Gang Liu, M. G. Rasul, and D. Harris. Techno-economic assessment of a residential grid-connected solar photovoltaic system. *Renewable Energy*, 53:85–94, 2013.

[13] Ping-Huan Kuo and Chiou-Jye Huang. An electricity price forecasting model by hybrid structured deep neural networks. *Sustainability*, 10(4):1280, 2018.

[14] Abdul Fatir Ansari, Caner Turkmen, Lorenzo Stella, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, et al. Chronos-Bolt: Efficient and accurate pretrained models for probabilistic time series forecasting. *arXiv preprint arXiv:2410.20081*, 2024.

[15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008, 2017.

[16] Rainer Storn and Kenneth Price. Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359, 1997.

[17] Nasimul Noman and Hitoshi Iba. Differential evolution for economic load dispatch problems. *Electric Power Systems Research*, 78(8):1322–1331, 2008.

[18] Makbul A. M. Ramli, H. R. E. H. Bouchekara, and Ali S. Alghamdi. Optimal sizing of PV/wind/diesel hybrid microgrid system using multi-objective self-adaptive differential evolution algorithm. *Renewable Energy*, 121:400–411, 2018.

[19] Jack Kiefer. Sequential minimax search for a maximum. *Proceedings of the American Mathematical Society*, 4(3):502–506, 1953.

[20] Francis X. Diebold and Roberto S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995.