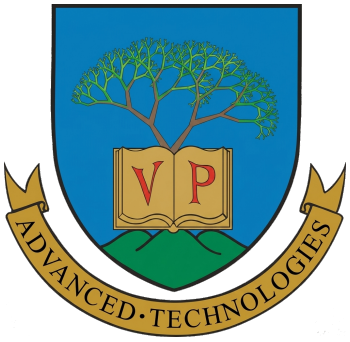
University of Pannonia
Faculty of Information Technology
Doctoral School of Information Science and Technology



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

PhD Thesis

Tamás Miseta

Supervisors:
Dr. Ágnes Vathy-Fogarassy

Dr. Attila Fodor

Field of science: Engineering Sciences

Branch of science: Information Sciences

Veszprém
2026

DOI:10.18136/PE.2026.1024

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Ágnes Vathy-Fogarassy, for her continuous guidance, encouragement, and invaluable support throughout the entire doctoral research process. Her expertise and dedication have been instrumental in shaping both the direction and the quality of this work.

I am equally grateful to Dr. Attila Fodor for his essential contributions as co-author, for sharing his deep knowledge of electrical energy systems, and for the many insightful discussions that strengthened the engineering aspects of this research.

I would like to thank Zsófia Závodi-Fodor for her collaboration on the photovoltaic energy system design and floating home energy system research, and for her practical expertise that enriched these contributions.

Finally, the financial support of the following institutions, grants, and programmes is gratefully acknowledged:

- The *New National Excellence Program of the Ministry of Human Capacities* (UNKP-20-2).
- *NVIDIA Corporation* for the donation of the Titan V GPU used for the neural network research.
- The *National Research, Development and Innovation Fund of Hungary*, financed under the K 19 funding scheme (project no. 131501).
- The *Ministry for Culture and Innovation* and the *Hungarian Research, Development and Innovation Fund* through the Cooperative Doctoral Programme (KDP 2023).

Abstract

Day-ahead electricity markets (DAMs) have been progressively liberalised and later reshaped by the integration of renewable energy sources. Photovoltaic (PV) and wind generation are non-dispatchable, and their growing share has increased price volatility. Energy storage can shift generation across hours, turning this volatility into an arbitrage opportunity. Over the past decade, falling PV and battery costs have strengthened the economic case. Machine learning methods for nonlinear price dynamics have also matured, and the energy sector is only beginning to apply them. Realising this opportunity requires three capabilities: forecasting electricity prices, optimising day-ahead trading strategies, and sizing the physical parameters of battery-enhanced PV energy systems for their intended operation. Contributions addressing each of these were developed and tested.

The first stage concerns day-ahead electricity price forecasting. At its core is the newly proposed Correlation-Driven Stopping Criterion (CDSC), a neural network regulariser that identifies the onset of overfitting from the rolling Pearson correlation between the training and validation losses and halts training accordingly. Across four benchmark problems and three network architectures, CDSC matched or exceeded conventional early stopping and reduced the test error by up to 2.89% relative to the full-training baseline, including on problems where early stopping brought little benefit. Separately, a hyperparameter-optimised univariate Long Short-Term Memory (LSTM) model was developed for the Hungarian Power Exchange (HUPX), reaching a mean absolute error (MAE) of 7.48 EUR/MWh. The model was then extended to the Germany EPEX SPOT, Singapore NEMS, and United Kingdom N2EX markets, where CDSC lowered the error in 10 of 12 configurations, by as much as 16.7% at the longer prediction horizon. Six pre-trained Transformer foundation models were also benchmarked under an optimised context window length. Evaluated in zero-shot setting, the best of them, most often Chronos 2, reached a lower MAE than the tuned LSTM in all six market-horizon scenarios, by 18.4–40.8%. This ranking held on the unseen 2025 holdout year.

The second stage optimises how a battery-enhanced PV plant trades on the 24-hour DAM once prices are forecast. A parametrised model of a DC-coupled hybrid inverter system defines the shared objective, and two newly proposed optimisation methods were developed for it. Optimised Energy Trading by Gradient Descent (OptEnT-GD) exploits the initial monotonicity of that objective. Evaluated across 42 system configurations, it raised daily profit by up to 46.26% over the immediate-dispatch baseline on the configurations with 5% storage loss, and under realistic forecasts its profit increase was only 1.93–6.93 percentage points below that of a perfect-information oracle. Optimised Energy Trading by Differential Evolution (OptEnT-DE) drops the monotonicity assumption and reaches a maximum profit about 3% lower than OptEnT-GD's at the largest battery capacity. The two methods are therefore complementary: OptEnT-GD is preferred when the initial monotonicity holds, and OptEnT-DE offers a derivative-free alternative for a broader class of objective functions. Across the configurations tested, the battery-capacity-to-inverter-power ratio emerged as the main driver of achievable profit.

The third stage sizes the physical system for its intended operation. For industrial prosumers barred from feeding surplus energy into the grid, the dichotomous search-based string optimisation (DSSO) algorithm, a consumption-aware method, and its evaluation framework were developed. For a single string, DSSO locates the optimum of any unimodal objective to within one panel of DC capacity in a logarithmic number of steps; multi-string systems are solved by coordinate-wise passes over the strings. In a case study on a Hungarian enterprise under a fixed tariff, DSSO reached roughly three times the self-reliance of a curtailment-capped baseline while lowering grid consumption by 14.8–17.5% and CO_2 emissions by 12.9–15%, at 37–45% of the conventional Sum Yearly Consumption Matching (SYCM) design size; under a DAM-based tariff it returned smaller systems still. For semi-autonomous structures with limited grid connectivity, a second methodology combines the Stochastic Consumption Estimation Algorithm (SCEA) for load synthesis, a coupled thermal-electrical model, and joint optimisation of the PV array and battery. In a year-long simulation-based case study of a real floating-home design, it reached 91.43% self-reliance and an 86.2% reduction in grid consumption. It also showed that oversizing the array beyond the weight-limited battery leaves 22.14 MWh of generation unutilised in a year.

Taken together, these methods span the forecasting, trading, and design stages of a battery-enhanced renewable power plant, and together they define the components of an end-to-end decision-support tool, whose coupled validation remains future work.

Kivonat

A másnapi villamosenergia-piacok (DAM-ek) fokozatos liberalizáción mentek keresztül, majd a megújuló energiaforrások integrációja alakította át őket. A fotovoltaikus (PV) és a szélenergia-termelés időjárásfüggő és nehezen szabályozható, növekvő részesedésük az energiamixben pedig fokozta a villamosenergia-árak volatilitását. Az energiatárolás lehetővé teszi a termelés órák közötti áthelyezését, így a volatilitás arbitrázslehetőséggé válik. Az elmúlt évtizedben a fotovoltaikus és az akkumulátoros rendszerek költségei csökkentek, ami javította a beruházások gazdaságosságát, a nemlineáris áramárdinamika kezelésére szolgáló gépi tanulási módszerek pedig kiforrottá váltak, energiaipari alkalmazásuk azonban még csak most kezdődik. E lehetőség kiaknázása három részfeladatot foglal magában: a villamosenergia-árak előrejelzését, a kereskedési stratégia optimalizálását és az akkumulátorral bővített fotovoltaikus rendszer tényleges működéshez igazított méretezését. Mindhárom részfeladathoz új módszertani megoldásokat dolgoztam ki és teszteltem.

A dolgozat első része a másnapi villamosenergia-árak előrejelzésével foglalkozik. Ennek középpontjában az újonnan javasolt Correlation-Driven Stopping Criterion (CDSC) áll: egy neurális hálózati regularizációs eljárás, amely a tanítási és a validációs veszteség közötti mozgóablakos Pearson-korreláció alapján ismeri fel a túltanulás kezdetét, és ennek megfelelően állítja le a tanítást. Négy referenciafeladaton és három hálózati architektúrán a CDSC elérte vagy felülmúlta a hagyományos early stoppingot, és a teszthibát a teljes tanítási referenciamodellhez képest legfeljebb 2,89%-kal csökkentette, azokon a feladatokon is, ahol az early stopping alig hozott javulást. Ettől függetlenül hiperparaméter-optimalizált, egyváltozós (univariate) Long Short-Term Memory (LSTM) modellt dolgoztam ki a magyar villamosenergia-tőzsde (HUPX) másnapi piacára, amely 7,48 EUR/MWh átlagos abszolút hibát (MAE) ért el. A modellt ezután kiterjesztettem a német EPEX SPOT, a szingapúri NEMS és az egyesült királysági N2EX piacra, ahol a CDSC a 12 konfigurációból 10-ben csökkentette a hibát, a hosszabb előrejelzési horizonton legfeljebb 16,7%-kal. Hat előtanított Transformer alapmodellt (foundation model) is kiértékeltem optimalizált kontextusablakhosszal. Feladatspecifikus továbbtanítás nélküli (zero-shot) beállításban a legjobb modell többnyire a Chronos 2 volt, mind a hat piac-horizont forgatókönyvben alacsonyabb MAE-t ért el, mint a paraméterhangolt LSTM (18,4–40,8%-kal). Ez a sorrend a 2025-ös, független tesztéven is fennmaradt.

A dolgozat második része a másnapi kereskedési stratégia optimalizálásával foglalkozik egy akkumulátorral bővített fotovoltaikus erőmű esetében a 24 órás DAM-on, miután a villamosenergia-árak előrejelzése megtörtént. A közös célfüggvényt egy paraméterezett, egyenáramú (DC) csatolású hibrid inverter modellje határozza meg, amelyre két, általam fejlesztett optimalizálási módszert alkalmaztam. Az Optimised Energy Trading by Gradient Descent (OptEnT-GD) kihasználja a célfüggvény kezdeti monotonitását. 42 rendszerkonfiguráción kiértékelve, az 5%-os tárolási veszteségű rendszereken a napi nyereséget az azonnali értékesítéshez képest maximum 46,26%-kal növelte, LSTM-alapú ár-előrejelzések mellett pedig a nyereségnövekedése 1,93–6,93 százalékponttal maradt el a tökéletes előrejelzést feltételező referenciáétól. Az Optimised Energy Trading by Differential Evolution (OptEnT-DE) elhagyja a monotonitási feltevést, és a legnagyobb akkumulátorkapacitásnál mintegy 3%-kal alacsonyabb maximális nyereséget ér el az OptEnT-GD-hez képest. A két módszer így kiegészíti egymást: az OptEnT-GD akkor előnyös, ha a kezdeti monotonitás fennáll, az OptEnT-DE pedig derivált nélküli alternatívát kínál a célfüggvények tágabb köréhez. A tesztelt konfigurációkban az akkumulátorkapacitás és az inverterteljesítmény aránya bizonyult az elérhető nyereség fő meghatározó tényezőjének.

A dolgozat harmadik része a fizikai rendszer tényleges működéshez igazított méretezésével foglalkozik. Azon ipari termelő-fogyasztók (prosumerek) számára, akik nem táplálhatnak be energiát a hálózatba, kifejlesztettem a fogyasztást is figyelembe vevő, dichotomous search-based string optimisation (DSSO) algoritmust és annak értékelési keretrendszerét. Egyetlen sztringre a DSSO bármely unimodális célfüggvény optimumát egyetlen panel egyenáramú teljesítményének pontosságáig, logaritmikus lépésszámban megtalálja; a többsztringes rendszereket a sztringeken végigfutó, koordinátánkénti menetek oldják meg. Egy magyar vállalatra vonatkozó, rögzített tarifás esettanulmányban a DSSO a leszabályozást (curtailment) korlátozó referenciamódszerhez képest mintegy háromszorosára növelte az önellátási arányt, miközben a hálózati vételezést 14,8–17,5%-kal, a CO_2-kibocsátást pedig 12,9–15%-kal csökkentette, a hagyományos Sum Yearly Consumption Matching (SYCM) tervezési módszerrel kapott méret 37–45%-án; DAM-alapú tarifa mellett ennél is kisebb rendszereket adott. A korlátozott hálózati csatlakozású, félig autonóm rendszerekhez egy második módszertant dolgoztam ki, amely a terhelési profil előállításához használt Stochastic Consumption Estimation Algorithm (SCEA) módszert, egy csatolt termikus-elektromos modellt és a fotovoltaikus rendszer és az akkumulátor közös optimalizálását egyesíti. Egy valós úszóház-terv egyéves, szimuláción alapuló esettanulmányában 91,43%-os önellátási arányt és 86,2%-os hálózativételezés-csökkenést ért el. A vizsgálat azt is megmutatta, hogy az akkumulátor súlykorlátozása miatt túlméretezett napelemes rendszer évente 22,14 MWh energiát hagy kihasználatlanul.

A javasolt módszerek együttesen lefedik egy akkumulátoros energiatárolóval bővített megújuló erőmű előrejelzési, kereskedési és tervezési szakaszait, és együtt egy teljes folyamatot átfogó döntéstámogató eszköz építőelemeit adják, amelynek összekapcsolt validációja további kutatás tárgya.

Zusammenfassung

Day-Ahead-Strommärkte (DAMs) wurden schrittweise liberalisiert und später durch die Integration erneuerbarer Energiequellen umgestaltet. Photovoltaik (PV) und Windkraft sind dargebotsabhängig und nicht steuerbar; ihr wachsender Anteil hat die Preisvolatilität erhöht. Energiespeicher können die Erzeugung zwischen den Stunden verschieben und diese Volatilität in eine Arbitrage-Möglichkeit verwandeln. Die im letzten Jahrzehnt stark gefallenen PV- und Batteriekosten haben die Wirtschaftlichkeit weiter gestärkt. Auch die maschinellen Lernmethoden für nichtlineare Preisdynamiken sind ausgereift, und der Energiesektor beginnt erst, sie anzuwenden. Die Nutzung dieser Möglichkeit umfasst drei Teilaufgaben: die Prognose der Strompreise, die Optimierung der Handelsstrategie und die Dimensionierung des batteriegestützten PV-Systems für den vorgesehenen Betrieb. Für jede dieser Stufen wurden neue methodische Lösungsansätze entwickelt und getestet.

Die erste Stufe betrifft die Day-Ahead-Strompreisprognose. In ihrem Zentrum steht das neu vorgeschlagene Correlation-Driven Stopping Criterion (CDSC), ein Regularisierungsverfahren für neuronale Netze, das den Beginn des Overfittings anhand der gleitenden Pearson-Korrelation zwischen Trainings- und Validierungsverlust erkennt und das Training entsprechend beendet. Über vier Benchmark-Probleme und drei Netzarchitekturen erreichte oder übertraf das CDSC das konventionelle Early Stopping und senkte den Testfehler gegenüber dem vollständig trainierten Referenzmodell um bis zu 2,89%, auch auf Problemen, bei denen Early Stopping kaum Verbesserung brachte. Unabhängig davon wurde ein hyperparameteroptimiertes univariates Long Short-Term Memory (LSTM)-Modell für die Ungarische Strombörse (HUPX) entwickelt, das einen mittleren absoluten Fehler (MAE) von 7,48 EUR/MWh erreichte. Das Modell wurde anschließend auf den deutschen EPEX-SPOT-, den Singapurer NEMS- und den N2EX-Markt des Vereinigten Königreichs ausgeweitet, wo das CDSC den Fehler in 10 von 12 Konfigurationen senkte, am längeren Prognosehorizont um bis zu 16,7%. Sechs vortrainierte Transformer-Foundation-Modelle wurden zudem mit optimierter Kontextfensterlänge untersucht. Im Zero-Shot-Betrieb erreichte das beste von ihnen, zumeist Chronos 2, in allen sechs Markt-Horizont-Szenarien einen niedrigeren MAE als das parameteroptimierte LSTM, um 18,4 bis 40,8%. Diese Rangfolge blieb im unabhängigen Testjahr 2025 bestehen.

Die zweite Stufe optimiert, wie ein batteriegestütztes PV-Kraftwerk auf dem 24-Stunden-DAM handelt, sobald die Preise prognostiziert sind. Ein parametrisiertes Modell eines DC-gekoppelten Hybridwechselrichter-Systems definiert die gemeinsame Zielfunktion, für die zwei neu vorgeschlagene Optimierungsmethoden entwickelt wurden. Optimised Energy Trading by Gradient Descent (OptEnT-GD) nutzt die anfängliche Monotonie dieser Zielfunktion. Über 42 Systemkonfigurationen steigerte es den Tagesgewinn bei den Konfigurationen mit 5% Speicherverlust um bis zu 46,26% gegenüber dem sofortigen Direktverkauf, und unter realistischen Prognosen lag die Gewinnsteigerung nur 1,93 bis 6,93 Prozentpunkte unter der eines Orakels mit perfekter Information. Optimised Energy Trading by Differential Evolution (OptEnT-DE) verzichtet auf die Monotonieannahme und erreicht bei der größten Batteriekapazität einen etwa 3% niedrigeren Maximalgewinn als OptEnT-GD. Die beiden Methoden ergänzen einander: OptEnT-GD ist vorzuziehen, wenn die anfängliche Monotonie gilt, während OptEnT-DE eine ableitungsfreie Alternative für eine breitere Klasse von Zielfunktionen bietet. Über die getesteten Konfigurationen hinweg erwies sich das Verhältnis von Batteriekapazität zu Wechselrichterleistung als wichtigster Treiber des erzielbaren Gewinns.

Die dritte Stufe dimensioniert das physische System für den vorgesehenen Betrieb. Für industrielle Prosumer mit Einspeiseverbot wurden der Algorithmus der dichotomous search-based string optimisation (DSSO), ein verbrauchsbewusstes Verfahren, und sein Bewertungsrahmen entwickelt. Für einen einzelnen String findet DSSO das Optimum jeder unimodalen Zielfunktion bis auf die Gleichstromleistung eines einzelnen Moduls in logarithmisch vielen Schritten; Mehrstring-Systeme werden durch koordinatenweise Durchläufe über die Strings gelöst. In einer Fallstudie an einem ungarischen Unternehmen erreichte DSSO bei fester Tarifstruktur etwa die dreifache Eigenversorgung einer abregelungsbegrenzten Vergleichsmethode, senkte den Netzbezug um 14,8 bis 17,5% und die CO_2-Emissionen um 12,9 bis 15%, bei 37 bis 45% der konventionellen Auslegung nach dem Sum Yearly Consumption Matching (SYCM); bei einem DAM-basierten Tarif fielen die Anlagen noch kleiner aus. Für halbautonome Strukturen mit eingeschränkter Netzanbindung kombiniert eine zweite Methodik den Stochastic Consumption Estimation Algorithm (SCEA) zur Lastprofilsynthese, ein gekoppeltes thermisch-elektrisches Modell und die gemeinsame Optimierung von PV-Feld und Batterie. In einer einjährigen simulationsgestützten Fallstudie eines realen Hausboot-Entwurfs erreichte sie 91,43% Eigenversorgung und eine Reduktion des Netzbezugs um 86,2%. Sie zeigte zudem, dass eine Überdimensionierung des PV-Felds gegenüber der gewichtsbegrenzten Batterie 22,14 MWh Erzeugung pro Jahr ungenutzt lässt.

Zusammengenommen decken diese Methoden die Prognose-, Handels- und Entwurfsstufen eines batteriegestützten erneuerbaren Kraftwerks ab und bilden gemeinsam die Bausteine eines durchgängigen Entscheidungsunterstützungswerkzeugs, dessen gekoppelte Validierung Gegenstand künftiger Arbeiten bleibt.

Contents

Abstract

Zusammenfassung

Abbreviations

Notations

List of Tables

List of Figures

List of Algorithms

1 Introduction **1**
1.1 Motivation and context . . . 1
1.2 Day-ahead electricity markets and price forecasting . . . 2
1.3 Energy trading strategies for battery-enhanced renewable power plants . . . 4
1.4 Physical design and optimisation of photovoltaic energy systems . . . 5
1.5 Research objectives and contributions . . . 7
1.6 Structure of the dissertation . . . 8

2 Day-Ahead Electricity Price Forecasting Using Machine Learning **10**
2.1 Introduction . . . 10
2.2 Correlation-Driven Stopping Criterion for neural network training . . . 12
2.2.1 Motivation and context . . . 12
2.2.2 Method . . . 13
2.2.3 Experimental methodology . . . 16
2.2.4 Results . . . 19
2.2.5 Discussion . . . 24
2.3 LSTM-based day-ahead electricity price forecasting with CDSC-enhanced training . . . 26
2.3.1 Motivation and context . . . 26
2.3.2 Method . . . 26
2.3.3 Experimental methodology . . . 26
2.3.4 Results . . . 30
2.3.5 Discussion . . . 35
2.4 Transformer-based foundation models for direct multi-step electricity price forecasting . . . 37
2.4.1 Motivation and context . . . 37

2.4.2 Method . . . 38
2.4.3 Experimental methodology . . . 39
2.4.4 Results . . . 39
2.4.5 Discussion . . . 45

3 Energy Trading Strategy Optimisation for Battery-Enhanced Power Plants 47
3.1 Introduction . . . 47
3.2 Parametrised system model . . . 48
3.2.1 Energy flow equations . . . 50
3.2.2 Physical constraints . . . 50
3.2.3 Objective function . . . 51
3.3 Adaptive gradient-descent-based energy trading optimisation . . . 52
3.3.1 Motivation and context . . . 52
3.3.2 Method . . . 52
3.3.3 Experimental methodology . . . 56
3.3.4 Results . . . 56
3.3.5 Discussion . . . 58
3.4 Differential-evolution-based energy trading optimisation and system sizing guide . . . 60
3.4.1 Motivation and context . . . 60
3.4.2 Method . . . 61
3.4.3 Experimental methodology . . . 63
3.4.4 Results . . . 63
3.4.5 Discussion . . . 64

4 Physical Design Optimisation of Photovoltaic Energy Systems 65
4.1 Introduction . . . 65
4.2 Consumption-aware dichotomous search-based string optimisation of multi-string photovoltaic systems . . . 67
4.2.1 Motivation and context . . . 67
4.2.2 Method . . . 68
4.2.3 Experimental methodology . . . 68
4.2.4 Results . . . 69
4.2.5 Discussion . . . 69
4.3 Design methodology for semi-autonomous electric energy systems . . . 72
4.3.1 Motivation and context . . . 72
4.3.2 Method . . . 72
4.3.3 Experimental methodology . . . 77
4.3.4 Results . . . 78
4.3.5 Discussion . . . 78

5 New Scientific Results 80
5.1 Thesis I: Day-ahead electricity price forecasting using machine learning . . . 80
5.2 Thesis II: Energy trading strategy optimisation for battery-enhanced power plants . . . 81
5.3 Thesis III: Physical design optimisation of photovoltaic energy systems . . . 82

6 Conclusion 83

Abbreviations

AC	Alternating Current
ARIMA	AutoRegressive Integrated Moving Average
BESS	Battery Energy Storage System
BIM	Building Information Modelling
BO	Bayesian Optimisation
CDSC	Correlation-Driven Stopping Criterion
CNN	Convolutional Neural Network
COP	Coefficient of Performance
CV	Cross-Validation
DAM	Day-Ahead Market
DC	Direct Current
DE	Differential Evolution
DM	Diebold–Mariano test
DSO	Distribution System Operator
DSSO	Dichotomous Search-based String Optimisation
ENTSO-E	European Network of Transmission System Operators for Electricity
EPF	Electricity Price Forecasting
EPS	Expanded Polystyrene
ESS	Energy Storage System
FFNN	Feed-Forward Neural Network
GARCH	Generalised AutoRegressive Conditional Heteroskedasticity
GBRT	Gradient-Boosted Regression Trees
GER	Germany Day-Ahead Market (EPEX SPOT)
GIS	Geographical Information System
HPC	High-Performance Computing
HPO	Hyperparameter Optimisation
HUPX	Hungarian Power Exchange
IPCC	Intergovernmental Panel on Climate Change
IRR	Internal Rate of Return
LCOE	Levelised Cost of Electricity
LEAR	Lasso Estimated AutoRegressive
LSTM	Long Short-Term Memory

MAE	Mean Absolute Error
MCS	Model Confidence Set
MILP	Mixed-Integer Linear Programming
MOPSO	Multi-Objective Particle Swarm Optimisation
MPC	Model Predictive Control
MPE	Mean Percentage Error
MPPT	Maximum Power Point Tracking
N-BEATS	Neural Basis Expansion Analysis for Time Series
NN	Neural Network
NPV	Net Present Value
OptEnT-DE	Optimised Energy Trading by Differential Evolution
OptEnT-GD	Optimised Energy Trading by Gradient Descent
PAC	Probably Approximately Correct
PSO	Particle Swarm Optimisation
PV	Photovoltaic
PVGIS	Photovoltaic Geographical Information System
RES	Renewable Energy Source
RF	Random Forest
RNN	Recurrent Neural Network
SCEA	Stochastic Consumption Estimation Algorithm
SP	Singapore Day-Ahead Market (NEMS)
SME	Small-to-Medium Enterprise
SoC	State of Charge
SYCM	Sum Yearly Consumption Matching
TPE	Tree-structured Parzen Estimator
TSO	Transmission System Operator
UK	United Kingdom Day-Ahead Market (N2EX)
UNECE	United Nations Economic Commission for Europe

Notations

C_{inside}	heat capacity of the interior
COP_{nom}	nominal coefficient of performance of the air-conditioning unit
T_{out}	outside (ambient) temperature
T_{target}	target inside temperature
T_{inside}	inside temperature
ΔT_{inside}	change in the inside temperature
CR	crossover rate of the differential evolution algorithm
C_{st}	battery storage capacity
D	number of stacked LSTM layers
d	discount rate
d_h	hidden dimension per LSTM layer
E_{grid}	energy consumed from the public grid
E_i^{PV}	energy generated by the PV system at hour i
E_{lost}	energy lost (curtailed) due to feed-in restriction
E_{st}	energy stored in the battery
$e_{\text{tr}}^{(j)}$	training error at epoch j
E_{used}	simulated energy consumption array
$e_{\text{va}}^{(j)}$	validation error at epoch j
$F(\vec{X})$	objective function for the trading optimisation
F_{de}	scale (mutation) factor of the differential evolution algorithm
$i_{\text{max}}^{\text{DE}}$	maximum number of iterations in the differential evolution algorithm
$i_{\text{max}}^{\text{GD}}$	maximum number of iterations in the gradient descent algorithm
k_{accop}	temperature-difference COP correction coefficient
k_{chem}	battery chemistry loss coefficient
k_{diss}	heat dissipation coefficient
k_{peb}	power electronics loss coefficient
k_{PV}	solar heating coefficient
L	lookback window size (timesteps of length Δt; hours on hourly markets)
L^*	optimal context length of a foundation model selected by grid search
$\vec{m}$	mutant individual in the differential evolution algorithm
N	maximum number of training epochs
m	epoch budget of the maximum number of epochs stopping criterion

m_{Best}	epoch budget of the selected best model
n	prediction horizon (timesteps of length Δt; hours on hourly markets)
n_{best}	epoch number of the selected best model
n_{stop}	epoch number when training was stopped
n_{μ}	epoch where the correlation first falls below μ_{cdsc}
$\text{obj}(P)$	objective function of the DSSO algorithm
$\mathcal{P}$	population in the differential evolution algorithm
P_{ac}	input power of the air-conditioning unit, $P_{\text{ac}} \geq 0$
s_{ac}	air-conditioning operating-mode indicator, +1 heating, −1 cooling
τ	self-reliance tolerance defining the near-optimal design set $\mathcal{D}_{\tau}$
P_{app}	appliance load power
P_{bin}	power injected to the battery
P_{bout}	power extracted from the battery
$P_{\text{bin,max}}$	maximum battery charging power (converter rating)
$P_{\text{bout,max}}$	maximum battery discharging power (converter rating)
p_{drop}	dropout rate
p_{es}	patience parameter of the early stopping criterion
$p_{\text{es,Best}}$	early stopping patience of the selected best model
P_{grid}	power consumed from the public grid
P_{max}	maximum output power of the inverter
$P_{\text{max}}^{\text{str}}$	maximum installed DC peak power (kWp) of a PV string
$P_{\text{max}}^{\text{app}}$	maximum power of an appliance
$P_{\text{min}}^{\text{str}}$	minimum installed DC peak power (kWp) of a PV string
$P_{\text{min}}^{\text{app}}$	minimum power of an appliance
$P_{\text{opt}}^{\text{str}}$	optimal installed DC peak power (kWp) of a PV string
P_{PV}	output power of the PV system
$P_{\text{PV,lost}}$	unutilised PV power
Q_{ac}	thermal power from the air-conditioning unit
Q_{app}	appliance heating power
Q_{diss}	heat dissipation transfer
Q_i	battery charge at hour i
Q_{inside}	total heat introduced inside the structure
Q_{max}	maximum battery capacity

Q_{vent}	ventilation heat loss
R	number of random-seed repetitions per trial
r	rolling Pearson correlation coefficient
s	population size in the differential evolution algorithm
$\vec{u}$	trial individual in the differential evolution algorithm
$\mathbf{W}_{\text{profile}}$	usage weight profile (24-element array of relative hourly usage weights)
$\vec{X}$	trading strategy vector, $\vec{X} = [x_1, x_2, \ldots, x_n]$
x_i	trading ratio in the i-th hour, in percent of inverter power, $x_i \in [-100, 100]$
$\vec{X}_{\text{opt}}$	optimised trading strategy vector
$\mathbf{X}_\Delta$	matrix of candidate trading strategies
Y	daily profit
Z_i	day-ahead market price at hour i
α	step size in the gradient descent method
$\alpha_{\min}$	minimum step size (exploitability threshold)
γ	number of variables simultaneously changed during gradient calculation
ΔE_{st}	change in stored energy
Δt	time resolution of the simulation (minutes)
ε	convergence tolerance of the DSSO algorithm
η_{st}	round-trip storage/inverter efficiency
Λ	fill ratio (fraction of time an appliance consumes high power while on)
λ_{cdsc}	patience parameter of the CDSC
$\lambda_{\text{cdsc,Best}}$	patience of the CDSC at the selected best model
ω_{Best}	window size of the CDSC at the selected best model
λ_{gd}	index of the row in $\mathbf{X}_\Delta$ with the highest profit increase
λ_{L1}	L1 (Lasso) regularisation penalty weight
λ_{L2}	L2 (Ridge) regularisation penalty weight
μ_{cdsc}	correlation threshold of the CDSC
$\mu_{\text{cdsc,Best}}$	correlation threshold of the CDSC at the selected best model
μ_{sr}	self-reliance ratio (percentage of time with no grid power required)
σ	step direction in the gradient descent method (+1: sell, −1: buy)
ϕ	variable storing the evaluated profit
ω	window size of the CDSC

List of Tables

2.1 NN architectures used for the CDSC benchmark evaluation. The four models cover the FFNN, CNN, and LSTM architectures, exercising the criterion across binary classification, multiclass classification, regression, and time-series tasks rather than a single architecture. 16

2.2 Parameter space for the CDSC grid search. Window size and patience are scanned in steps of 5, whereas the correlation threshold is swept across its full range $[-1, 1]$ at a finer resolution (step 0.01), the parameter to which the stopping decision proves most sensitive. 18

2.3 Metrics for the maximum number of epochs criterion with the best parameter m_{Best}. Capping the epoch budget lowered the test error on all four datasets, most on MNIST (−2.73%) and only marginally on Boston Housing and Gold. 20

2.4 Metrics for the early stopping criterion with the best patience parameter $p_{\text{es,Best}}$. The largest gain was on Credit Card Fraud Detection (−1.25%), while Boston Housing saw no improvement at any patience. 21

2.5 Metrics for the CDSC with the best parameter combination ($\mu_{\text{cdsc,Best}}$, ω_{Best}, $\lambda_{\text{cdsc,Best}}$) for each dataset. The CDSC lowered the test error on every dataset, most on MNIST (−2.89%); the best threshold is positive for the cleaner Credit Card Fraud Detection and MNIST problems and negative for the smaller, noisier Boston Housing and Gold problems. 21

2.6 Mean percentage error change relative to baseline for the three stopping criteria with their best parameter settings. Negative values indicate improved out-of-sample performance. 22

2.7 Parameter importance values of the CDSC, extracted from fitted random forest models, together with test accuracies of the random forest classifiers. The correlation threshold μ_{cdsc} dominates on every dataset (importance ≥ 0.70), whereas the window size and patience contribute comparatively little. 23

2.8 BO (GBRT) convergence statistics: number of evaluations needed to find a sufficiently good CDSC parameter setting (50 runs per problem). BO located a competitive setting within a small fraction of the search space (a median of 18.5–37 evaluations against the 1800–4300 combinations that the parameter ranges yield when the threshold is discretised at 0.05 intervals); the median falls well below the mean because a few unfavourable initialisations lengthen the tail. 23

2.9 Hyperparameters and their tested values for the LSTM grid search on HUPX. The second-phase grid pairs each lookback window and layer width with each regularisation method, holding the linear activation and $[0, 1]$ normalisation fixed from the first phase. 27

2.10 Descriptive statistics of the GER price series (EUR/MWh). The mean price rose from 30.4 in 2020 to 236.5 at the 2022 energy-crisis peak before easing to 77.8 in 2024, and negative prices occur in every year. 28

2.11 Descriptive statistics of the SP price series (SGD/MWh). Prices are far more volatile than on the European markets, with the mean peaking at 292.9 in 2022 and half-hourly maxima between 4477 and 4545 in every year from 2021. 28

2.12 Descriptive statistics of the UK price series (GBP/MWh). The mean peaked at 240.5 in 2022 and fell to 85.8 by 2024, with renewable-driven negative prices in every year. 28

2.13 Mean absolute errors (EUR/MWh) for different activation functions and normalisation range combinations. The lowest error (MAE = 8.03) was obtained with linear activation and $[0, 1]$ normalisation. 30

2.14 Effect of regularisation methods on the best-performing LSTM network topology (168 LSTM neurons, 48-hour lookback, linear activation, $[0, 1]$ normalisation). The unregularised network is the most accurate (MAE = 7.48 EUR/MWh); every explicit weight or dropout penalty raises the error, indicating that the LSTM gates already provide sufficient regularisation on this task. 30

2.15 LSTM results across datasets, prediction horizons, CV types, and stopping methods. All MAE values are computed on the held-out 2024 test set. Results are ordered by test MAE within each scenario (best first). The CDSC attained a lower test MAE than early stopping in 10 of the 12 matched configurations, with the advantage concentrated at the longer prediction horizon. 31

2.16 Optimised stopping criterion parameters for LSTM models across all configurations. The CDSC settings vary widely across markets and horizons (threshold μ_{cdsc} from 0.0 to 0.7, window from 10 to 20, patience from 5 to 30), which is why they were tuned jointly with the architecture rather than fixed in advance. 33

2.17 LSTM stopping-criterion comparison (CDSC versus early stopping) by a two-sided DM test with the Harvey–Leybourne–Newbold small-sample correction, at the 5% significance level. Δ MAE is the CDSC test MAE minus the early-stopping test MAE; negative values indicate a lower error for the CDSC. "Sig." marks whether the difference is statistically significant. Results are shown for the 2024 test year and the 2025 holdout year. 34

2.18 LSTM test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; the year 2025 was not seen during tuning. **Bold** marks the lower MAE per column within the LSTM family. The relative ranking of the two stopping methods largely carries over from 2024 to the unseen 2025 year. 35

2.19 Membership of the LSTM models in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Neither stopping criterion entered the set in any scenario. 35

2.20 Transformer-based foundation models evaluated in zero-shot inference. Model sizes are reported as in the original publications. The six models are drawn from three families (Chronos, TimesFM, Moirai) and span 9M to 205M parameters, all applied without any task-specific gradient updates. . 38

2.21 Foundation model results across datasets and prediction horizons. Models are ordered by test MAE within each scenario. Chronos 2 attained the lowest test MAE in five of the six scenarios, the exception being UK at $n = 40$, where Moirai-2 (Small) was marginally lower. 40

2.22 Comparison of the best foundation model (FM) with the best LSTM configuration in each scenario. The best FM is Chronos 2 in every scenario except UK at $n = 40$, where Moirai-2 (Small) is marginally lower. CV: cross-validation; ES: early stopping. The improvement is computed as $(\mathrm{MAE}_{\mathrm{LSTM}} - \mathrm{MAE}_{\mathrm{FM}})/\mathrm{MAE}_{\mathrm{LSTM}} \times 100\%$. 41

2.23 Best conventional baseline configurations (test MAE on the 2024 test set) across the six scenarios. For LEAR and N-BEATS the better of the single-fold and CV configurations is shown; the seasonal-naive baseline has no tunable parameters. L: lookback window (timesteps). N-BEATS was the most accurate baseline in five of the six scenarios, yet even the strongest baseline remained less accurate than the best foundation model in every scenario. 42

2.24 Conventional baseline test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; 2025 was not seen during tuning. **Bold** marks the lower error of LEAR and N-BEATS within each scenario-year column. N-BEATS leads on SP and UK in both years, and the seasonal-naive floor trails throughout. 42

2.25 Membership of the conventional baselines in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Only N-BEATS entered, and only in two scenarios. 43

2.26 Regime-segmented test MAE and directional accuracy on the 2024 test year at the short horizon ($n = 4$). Spike: price above the annual mean plus two standard deviations; negative: price below zero; normal: the remainder. Directional accuracy (Dir. acc.) is the fraction of within-horizon move directions predicted with the correct sign. **Bold** marks the best value per column within a market; the SP negative-price column is left unmarked because that regime, with only 12 points, is not interpretable. 43

2.27 Foundation-model test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; the year 2025 was not seen during tuning. **Bold** marks the lowest MAE per column within the family. Chronos 2 remains the most accurate model in almost every column across both years, confirming that its lead is not specific to the 2024 test year. 44

2.28 Membership of the foundation models in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Chronos 2 entered in every scenario and was the sole member on GER at both horizons and on UK at $n = 4$. 44

3.1 Variables and their values evaluated across all scenarios. All possible combinations of system parameters were tested with each trading strategy and optimisation method. The three inverter powers, seven battery capacities, and two storage efficiencies yield the 42 system configurations used throughout the chapter. 56

3.2 Results of OptEnT-GD in the oracle scenario, grouped by minimum step size. *No. of best results* indicates the number of configurations where the step size achieved the highest profit. Cross-tabulation values show the average relative profit difference between step sizes compared to the immediate-dispatch baseline. The step size $\alpha_{\min} = 5$ won the most configurations (14), whereas $\alpha_{\min} = 25$ gave the highest average profit across the 42 configurations. 57

3.3 Results of OptEnT-GD in the predicted scenario, grouped by minimum step size. The optimal step size shifted from $\alpha_{\min} = 25$ in the oracle scenario to $\alpha_{\min} = 50$ in the predicted scenario. 57

3.4 Comparison of OptEnT-DE (100 runs) and OptEnT-GD with predicted day-ahead prices, $Q_{\max} = 12$ kWh, $\eta_{\text{st}} = 0.95$. OptEnT-GD yields about 3% higher profit than the best OptEnT-DE run in this large-battery case. 63

4.1 System sizes (in kWp) designed with the SYCM method, the E_{lost} limit method, and the proposed DSSO method for each surface configuration. East-west configurations show string 1 + string 2 sizes. Under the fixed tariff the DSSO systems are sized at 37–45% of the corresponding SYCM systems, and the DAM tariff yields smaller systems still. 70

4.2 Yearly aggregated energy metrics for systems designed with the SYCM, E_{lost} limit, and DSSO methods across the four surface configurations. E_{grid}: grid consumption (MWh); E_{lost}: curtailed energy (MWh); μ_{sr}: self-reliance ratio (%); CO_2: yearly emissions (t). Relative to the E_{lost} limit method, the fixed-tariff DSSO systems roughly triple the self-reliance ratio and lower both grid consumption and CO_2 emissions. 71

4.3 Percentage decrease in grid consumption for each PV-storage combination. All values are percentages. The full PV system with the largest storage (20 kWh) gives the greatest reduction (86.2%), and the reduction grows with storage size for every configuration. 78

4.4 Self-reliance ratio μ_{sr} (%) for each PV-storage combination. Self-reliance peaks at 91.43% for the full PV system with 20 kWh storage, with the roof-only system reaching almost the same value (89.59%). 78

4.5 Unutilised (curtailed) PV energy in MWh for each PV-storage combination. The full PV system leaves the most energy unutilised (22.14 MWh with the largest battery), and the amount falls as installed PV capacity is reduced, exposing the oversizing penalty under weight-limited storage. 78

List of Figures

1.1 Structure description of the dissertation: the prediction–control–design decision-support pipeline for battery-enhanced renewable power plants. Thesis I forecasts day-ahead prices, Thesis II optimises the trading strategy from those forecasts, and Thesis III designs the physical battery-enhanced PV energy system on which the strategy operates. 9

2.1 Visual representation of the Correlation-Driven Stopping Criterion. The rolling Pearson correlation coefficient of the training and the validation error is calculated for ω-wide windows step-by-step. n_μ denotes the epoch where the rolling correlation coefficient falls first below the threshold μ_{cdsc}. The training stops at epoch n_{stop}, and n_{best} denotes the epoch of the best validation error selected. 14

2.2 Graphical representation of the comparative evaluation methodology. For each of the 50 shuffled datasets, models are trained for N epochs, producing complete epoch logs. The three stopping criteria are simulated on these shared logs by truncation, ensuring identical training conditions across all comparisons. 17

2.3 Effect of the maximum epoch parameter m on the performance of the maximum number of epochs stopping criterion. Top panels: mean percentage error change ($\overline{\%e}_{\text{ts}}^{ME(m)}$) as a function of m. Bottom panels: average stopping epoch ($\overline{n_{\text{stop}}}^{ME(m)}$) and average selected epoch ($\overline{n_{\text{best}}}^{ME(m)}$). An intermediate m lowers the test error below the full-budget baseline on all four datasets, confirming that the epoch budget itself acts as a regularising hyperparameter. 19

2.4 Effect of the patience parameter p_{es} on the performance of the early stopping criterion. Top panels: mean percentage error change ($\overline{\%e}_{\text{ts}}^{ES(p_{\text{es}})}$) as a function of p_{es}. Bottom panels: average stopping and selected epochs. Early stopping reduced the test error over a wide patience range on MNIST and Credit Card Fraud Detection, but no patience value improved on the baseline for Boston Housing, where the validation loss plateaus. 20

2.5 Regret bound analysis of the three stopping criteria across the four benchmark problems. For each criterion, the regret distribution of the 50 best-performing parameter settings is shown. The CDSC achieves lower mean regret and lower variance than both the maximum number of epochs and early stopping criteria in all cases. 22

2.6 Best achievable $\overline{\%e}_{\text{ts}}$ with optimal μ_{cdsc} for each fixed λ_{cdsc}–ω combination. Darker blue regions indicate parameter regions where the CDSC achieves the greatest improvement. A connected trench of favourable ω–λ_{cdsc} combinations is visible for Credit Card Fraud Detection, MNIST, and Gold, but is far less pronounced for Boston Housing. 24

2.7 Normalisation ranges tested jointly with the three activation functions (linear, sigmoid, tanh). The bounds of the normalisation range were designed to match the output range of each activation function. 27

2.8 The two temporal validation strategies and the two held-out evaluation years. Panel (a): single-fold validation, with 2020–2021 and the first part of 2022 as the training set and the remainder of 2022 together with 2023 as the validation set. Panel (b): four-fold leave-one-year-out CV over 2020–2023, each year serving once as the validation fold while the remaining three form the training set. Hatched regions mark the embargo zones of width $L + n$ inserted at every partition boundary to prevent leakage from lagged features. Both the 2024 test year and the 2025 robustness year sit outside every tuning and model-selection step, which is what makes the out-of-year comparison of Section 2.3.4 a genuine holdout. 29

3.1 Structure of the proposed battery-enhanced PV energy system. The hybrid inverter manages energy flow among the PV array, BESS, the public grid, and the energy exchange. The EPF module provides forecast inputs to the trading strategy optimisation. 49
3.2 Absolute (a) and percentage (b) profit increase of OptEnT-GD in the oracle scenario. 58
3.3 Absolute (a) and percentage (b) profit of OptEnT-GD in the predicted scenario. 58
3.4 Percentage profit difference between the predicted and oracle scenarios for different system configurations (PV = $P_{\max}$). Across the 42 configurations, the prediction error lowers the profit increase over the immediate-dispatch baseline by only 1.93 to 6.93 percentage points. 59

4.1 The recommended workflow for the DSSO algorithm. Surface data and consumption data are combined with reference PV generation data from the PVGIS tool to form the dataset, which is then optimised by the DSSO algorithm using economic metrics to produce the optimal string configuration. 69
4.2 Block diagram of the simulation workflow for the semi-autonomous energy system design methodology. The SCEA generates load profiles from appliance properties, which are combined with PV output data, weather data, and electricity prices to form the input dataset for the system model. . . . 73

List of Algorithms

1 CDSC: Correlation-Driven Stopping Criterion. Training stops once the rolling correlation between the training and validation losses has been below the threshold μ_{cdsc} for a cumulative total of λ_{cdsc} epochs, and the epoch with the best validation error (n_{best}) is returned. 15

2 OptEnT-GD: Adaptive gradient-descent-based energy trading optimisation. Each iteration selects the most profitable multi-hour transaction and exploits it with an adaptive step size, stopping when no transaction improves the daily profit. 53

3 GradCalc: Gradient calculation for multi-hour transactions. All n^{γ} candidate transactions are evaluated and the index of the one with the largest profit increase is returned, or -1 if none improves the current strategy. 54

4 AdaptiveStep: Adaptive step-size calculation. Exploiting the initial monotonicity of the profit function, an accelerating search first overshoots the optimum and a refinement phase then locates the local maximum nearest zero. 55

5 OptEnT-DE: Differential-evolution-based energy trading optimisation. A population of 24-hour trading schedules is evolved through mutation, crossover, and greedy selection, and the fittest schedule in the final population is returned. 62

6 The DSSO algorithm for single-string PV sizing. Dichotomous (bisection) search maximises the unimodal objective function $\text{obj}(P)$ by discarding half of the power interval at each step, converging in $O(\log_2((P^{\text{str}}_{\text{max}} - P^{\text{str}}_{\text{min}})/\varepsilon))$ iterations to the optimal installed peak power $P^{\text{str}}_{\text{opt}}$ of the string. 70

7 The SCEA for synthetic load-profile generation. For each time slot a random draw is scaled by the appliance's hourly usage weight, fill ratio, and power range; the product of the draw and the fill ratio acts as a continuous usage-intensity factor rather than an on/off state. The procedure returns one day's profile for one appliance; it is invoked independently for each of the 365 days and each appliance, and the results are summed slot-wise to give the year-long load profile at resolution Δt. 74

Chapter 1
Introduction

1.1 Motivation and context

The global transition toward weather-dependent renewable energy sources (RES) has introduced complex and interconnected challenges across the electricity value chain [1]. According to the Renewable Energy Directive of the European Union (2018/2001/EU, amended by 2023/2413/EU), the share of renewable energy in gross final energy consumption must reach at least 42.5% by 2030, with a target of 45% [2]. Among renewables, solar and wind energy account for the largest share of new capacity additions [3]; however, their inherent intermittency and limited dispatchability pose significant challenges for Transmission System Operators (TSOs) and Distribution System Operators (DSOs), who must maintain a continuous balance between generation and demand to ensure grid stability [1]. The progressive liberalisation of electricity markets since the 1990s has reshaped this balancing problem into a price-based one: in modern day-ahead markets (DAMs), generators, retailers, and large consumers submit price-quantity bids for the delivery blocks of the following day, and the resulting clearing prices reflect short-term scarcity [1]. As the share of non-dispatchable wind and solar generation has grown, these prices have become increasingly volatile through the merit-order effect, first quantified for wind power in Germany by Ketterer [4] and later extended to photovoltaic (PV) generation by Cludius et al. [5]. This volatility creates both risks and opportunities for renewable producers: coupling weather-dependent generation with energy storage allows a power plant to shift output over hours, decoupling sales from instantaneous resource availability and selling energy when prices are highest [6]. These energy systems, which are the main subject of this dissertation, are referred to as battery-enhanced renewable power plants.

Three interconnected research challenges arise from this context. First, the profitable operation of battery-enhanced renewable power plants in DAMs requires accurate electricity price forecasting (EPF) [1], as forecast errors translate directly into financial losses. Second, given reliable price forecasts, the energy trading strategy of buying, storing, and selling energy must be optimised to maximise profit subject to the physical constraints of the battery energy storage system (BESS) [6]. Third, the physical design of the PV energy system itself must be optimised under real-world regulatory, economic, and sustainability constraints that vary across installation contexts, from industrial prosumers prohibited from grid injection [7] to semi-autonomous structures with limited grid connectivity.

These three challenges form a coherent research pipeline: prediction methods serve as an input to the trading strategy optimisation, the trading strategy motivates the need for battery-enhanced physical design, and the physical design constrains the system that ultimately trades on the market. All three pillars are addressed through a series of interrelated contributions spanning machine learning, mathematical optimisation, and energy systems engineering.

1.2 Day-ahead electricity markets and price forecasting

In competitive DAMs, market participants, including generators, retailers, and large industrial consumers, submit price-quantity bids before physical delivery, making accurate EPF essential for financially rational decision-making [1]. Forecast errors can directly lead to financial losses through suboptimal bidding strategies, over- or under-contracting of forward positions, and inefficient dispatch scheduling. Beyond individual market participants, reliable price signals also assist grid operators with congestion management and facilitate the integration of price-responsive demand flexibility [8], which is increasingly crucial as the proportion of non-dispatchable renewable generation expands.

Accurate day-ahead EPF has been a central research problem since the deregulation of power markets in the 1990s [1]. Electricity prices exhibit several characteristics that make forecasting particularly challenging, including strong multi-scale seasonality (intraday, weekly, and annual cycles), non-stationarity driven by fuel price shocks and evolving generation mixes, and heavy-tailed distributions featuring extreme spikes and occasional negative prices caused by renewable surplus. These stylised facts were catalogued in Weron's statistical monograph on power-market modelling [9] and reassessed a decade later in his state-of-the-art review of EPF methodology [1]. These characteristics also vary considerably across markets. Interconnected European markets such as the Germany DAM, operated by the European Power Exchange (EPEX SPOT), typically exhibit moderate volatility, as documented in the cross-market review of Lago et al. [8], which compares five European bidding zones under a common protocol. More isolated systems, such as the Singapore DAM, operated as the National Electricity Market of Singapore (NEMS), experience more frequent extreme price events; this pattern is examined by Loi and Jindal [10] in their post-deregulation assessment of Singapore wholesale prices in *Energy Policy*. The large-scale integration of variable RES has further increased price volatility and disrupted traditional seasonal patterns through the merit-order effect. Ketterer [4] quantified the wind-driven component of this effect for the Germany DAM, reporting a statistically significant downward shift in mean prices alongside higher conditional volatility as wind penetration rises. Cludius et al. [5] extended the analysis to the 2008–2016 period and to PV generation, estimating the joint merit-order effect of wind and PV in Germany together with its distributional consequences for consumers and producers. Consequently, a forecasting methodology must remain robust across diverse market structures and operating conditions [8].

Early EPF approaches relied on statistical models such as AutoRegressive Integrated Moving Average (ARIMA) [11], regime-switching, and Generalised AutoRegressive Conditional Heteroskedasticity (GARCH) [12] variants, which provide interpretable baselines but are limited in their ability to capture nonlinear dynamics. Contreras et al. [13] introduced an ARIMA formulation for next-day prices in the mainland Spanish and Californian pools, establishing an early linear benchmark for the problem. Mount et al. [14] addressed the spike regime with a regime-switching model with time-varying parameters, explicitly targeting jump behaviour that linear ARIMA cannot capture. Garcia et al. [15] concentrated on the conditional variance structure, deriving a GARCH forecaster reported to outperform a general ARIMA on volatile Spanish and Californian day-ahead series. Wavelet-ARIMA hybrids [16] and jump-diffusion models [17] have extended the statistical toolkit, yet the restrictive parametric form these models impose on the price process constrains their applicability to markets with increasingly complex price-formation mechanisms.

Machine learning methods, including support vector regression [18], random forests [19], and gradient boosting frameworks such as eXtreme Gradient Boosting (XGBoost) [20] and Light Gradient Boosting Machine (LightGBM) [21], have achieved state-of-the-art results in multiple markets by learning complex nonlinear relationships directly from data. The open-access benchmark of Lago et al. [8] places the best tuned deep and shallow models ahead of statistical baselines across five European DAMs and provides the reference protocol adopted in later EPF studies. Lehna, Scheller, and Herwartz [22] compared time-series

and neural network (NN) models for the Germany DAM while explicitly incorporating external regressors, reporting that gradient-boosted and Long Short-Term Memory (LSTM) pipelines augmented with fundamentals dominate univariate econometric baselines. Feature engineering, particularly the selection of lag window depth and calendar encoding, plays a critical role in the performance of these models. Uniejewski et al. [23] introduced automated variable selection and shrinkage for day-ahead EPF that subsequently became known as the Lasso Estimated AutoRegressive (LEAR) benchmark, and Lago et al. [8] adopted it as the canonical statistical reference in their cross-market benchmarking study, which also highlighted the lack of rigorous evaluation standards in EPF and advocated for public datasets, the Diebold–Mariano test [24], and standardised evaluation protocols.

Deep learning architectures have further advanced EPF. The LSTM network [25] addresses the vanishing gradient problem of vanilla recurrent neural networks (RNNs) through its gating mechanism, and has been widely applied to EPF. Ugurlu et al. [26] applied RNNs, including unidirectional and bidirectional LSTM configurations, to day-ahead EPF. Kuo and Huang [27] proposed a hybrid structured deep NN for EPF that combines convolutional and LSTM components; LSTM has also been used for related load-forecasting tasks [28]. The direct multi-step forecasting strategy, which assigns a separate output dimension to each prediction horizon, avoids the error compounding inherent in recursive approaches. Ben Taieb et al. [29] compared recursive, direct, and hybrid strategies across the 111 series of the NN5 forecasting competition and characterised when each strategy is preferable. Marcjasz et al. [30] carried this comparison into the EPF setting by contrasting single-output (recursive) and multiple-output (direct) NN architectures for day-ahead EPF and reporting the multiple-output formulation as the more accurate of the two on European DAMs. However, LSTM-based models for EPF are typically deployed with manually chosen lag window lengths and fixed early stopping configurations: Prechelt's catalogue of patience-based early stopping variants [31] remains the de facto regularisation recipe, and the EPF benchmarking review of Lago et al. [8] reports that hyperparameter optimisation in the field most often fixes feature-engineering choices to hand-picked values. As a result, LSTM models have not been subjected to large-scale joint hyperparameter optimisation that includes these design choices.

Transformer-based architectures, including the Informer [32], Autoformer [33], Temporal Fusion Transformer [34], and Patch Time Series Transformer (PatchTST) [35], have introduced attention-based mechanisms that model long-range dependencies without recurrence. However, their advantage over simple models can diminish on moderate-sized datasets [36]. Most recently, pre-trained time-series foundation models have emerged as a new paradigm. The original Chronos of Ansari et al. [37] tokenises continuous values into a fixed integer vocabulary and applies language-model-style cross-entropy training to a Transformer backbone; its successor Chronos-Bolt [38] replaces the per-token autoregressive head with a direct patch-based decoder for faster inference at comparable accuracy. TimesFM [39] uses a patched decoder-only architecture pre-trained on a large corpus of real-world and synthetic time series, operating directly on continuous values. Moirai [40] employs a masked-encoder Transformer with frequency-aware patch embeddings and a mixture-of-distributions output head, trained on a heterogeneous multi-domain time-series corpus. These models enable zero-shot forecasting without task-specific training [41]. Early benchmarks on EPF tasks indicate that foundation models perform on par with traditional approaches but do not yet consistently outperform well-tuned task-specific models [41]. Notably, the sensitivity of these foundation models to the context window length has not been systematically examined on DAM datasets.

Beyond model architecture, the training methodology itself is a critical and often under-explored component [8]. The stopping criterion determines when training terminates and directly affects generalisation performance. Classical early stopping monitors validation loss and halts training when it fails to improve for a specified patience period [31]; however, it monitors only a single loss curve, and its hyperparameters are not tuned systematically in many cases. Hyperparameter optimisation (HPO) in EPF is applied inconsistently: most studies tune only model-architecture parameters and leave feature-engineering choices fixed [8]. Cross-validation (CV) for time-series data requires temporal awareness to prevent information leakage. Tashman [42] reviewed out-of-sample tests of forecasting accuracy and laid out the rolling-origin evaluation protocol that replaces random folds with chronologically ordered forecasts, establishing the common ground rules for time-series error estimation. Bergmeir and Benítez [43] later examined the question for stationary autoregressive series and reported that blocked and last-block CV are appropriate in that setting, whereas independent-sample k-fold CV is unreliable when temporal dependence is non-trivial. Standard k-fold CV violates the temporal ordering assumption, while holdout, expanding window, and sliding window approaches preserve it at the cost of fewer evaluation points [44]. When lag-based features span partition boundaries, purged and embargoed CV, which removes observations near fold boundaries, is necessary to eliminate look-ahead contamination [45]. The systematic comparison of CV strategies for assessing the quality of HPO in EPF remains an open question [8].

Despite these advances, several methodological gaps persist. First, early stopping monitors only the validation loss, which can be suboptimal. The joint movement of the training and validation losses could be examined together, which could be a basis for a novel stopping criterion. Second, existing LSTM models combined with such a criterion have not been tested with comprehensive hyperparameter optimisation that jointly tunes the model architecture, lag window length, CV type, and stopping-criterion parameters [8]. Third, although there is preliminary research on time-series foundation models [41], their context-length sensitivity remains unexplored in EPF.

1.3 Energy trading strategies for battery-enhanced renewable power plants

The economic viability of battery-enhanced renewable power plants depends critically on the energy trading strategy employed on the DAM [6]. BESSs allow power plant operators to decouple the timing of energy generation from the timing of sale, enabling them to buy energy when prices are low, store it, and sell it when prices are high. This temporal arbitrage can significantly increase profit compared to the immediate-dispatch baseline, where all generated energy is sold immediately at the prevailing market price [6].

Two main types of battery-enhanced PV power plants (PVPPs) exist: AC-coupled and DC-coupled systems [46]. Both have been studied extensively, primarily from the control and optimisation perspective. In DC-coupled systems, the PV array, battery, and inverter share a common DC bus, which simplifies the system model and reduces conversion losses [46]. Hybrid grid-tied solar inverters with a BESS, which combine on-grid and off-grid properties [47], form the basis of the commercially available systems considered in this work.

Control strategies for battery-enhanced power plants have evolved from simple to increasingly sophisticated approaches. Early controllers were rule-based: Garcia et al. [48] presented two-hourly energy supervisory controls for off-grid hybrid systems, and Dash and Bajpai [49] addressed the frequent start-stop cycling of fuel cells. Cozzolino et al. [50] proposed a power management strategy for a hybrid PV energy system on artificial islands. Fuzzy logic [51] enabled more elaborate rulesets but lacked predictive capabilities. With the adoption of artificial NNs [52], solutions became increasingly data-driven and generalisable.

Optimisation methods have been applied extensively to energy trading and microgrid control. Dynamic programming was utilised for microgrid control [53]. Khalid et al. [6] maximised the profit of wind-battery stations using price-forecasting-based control. Meta-heuristic approaches, particularly Particle Swarm Optimisation (PSO), were used directly to develop control strategies [54]. Baghaee et al. [55] applied multi-objective PSO to a stand-alone wind/PV/fuel-cell microgrid to obtain Pareto-optimal designs balancing reliability and cost. The same authors subsequently extended this line of work in [56], applying multi-objective PSO (MOPSO) to the joint power-management and sizing of a wind/PV microgrid with hydrogen energy storage. Kaviani et al. [57] described a PSO-based sizing method driven by profitability. Other approaches include sliding mode control [58], mixed-integer linear programming (MILP) for combined heating and cooling power systems [59], adaptive search algorithms for wind power plants [60], and Markowitz theory-based risk-aware control [61].

Several model predictive control (MPC) solutions have been published for energy systems. MPC has been applied to smoothing the power output of wind power plants [62], frequency control [63], microgrid control [64], and combined cooling, heating, and power systems [59]. Zhuo et al. [65] described an MPC approach that merges profit optimisation with battery life optimisation. In contrast, Chen et al. [66] proposed a strategy combining LSTM prediction with Deep Deterministic Policy Gradient decision-making. A common characteristic of these MPC-based systems is that they introduce a new set of prediction problems; the quality of the control strategy is fundamentally limited by the accuracy of the price forecast [67].

Heydari et al. [67] demonstrated that EPF is key for profitable trading strategies in isolated power grids. Despite this comprehensive body of literature, complete model predictive control workflows that integrate accurate EPF with trading strategy optimisation for parametrised system models remain scarce.

Several gaps remain in the trading strategy literature. First, although there are numerous BESS models in the literature [46], a custom simplified and parametrised system model is required that captures the DC-coupled constraints governing hourly trading (battery capacity, storage efficiency, and inverter power limits) as a continuous objective for hourly trading optimisation. Second, no adaptive gradient-based method has been developed specifically for the energy trading control domain. Third, differential evolution (DE) [68], despite its successful application to economic dispatch [69] and renewable energy system sizing [70], has not been adapted to the energy trading domain, which has a specific constraint structure.

1.4 Physical design and optimisation of photovoltaic energy systems

The physical design of a PV energy system, including panel configuration, string sizing, orientation, storage capacity, and inverter rating, directly determines the system's energy yield, economic return, and environmental impact. Liu et al. [71] demonstrated this coupling using a techno-economic simulation and optimisation method for residential grid-connected PV that jointly tracks energy yield, economic metrics, and carbon dioxide emissions. Fan and Xia [72] similarly tied design choices to economic and environmental outcomes by formulating PV retrofit of an apartment complex as a multi-objective optimisation problem driven by net present value (NPV) and payback period. Industrial and household-size prosumers face diverse constraints that render conventional design methods suboptimal, including regulatory restrictions on grid injection, physical limitations of installation sites, and increasingly important corporate sustainability targets [7].

The most common PV energy system design approach matches a system's annual sum of output to the site's annual sum of consumption (Sum Yearly Consumption Matching, SYCM). The techno-economic study of Liu et al. [71] for residential grid-connected PV exemplifies

this design philosophy when feed-in is permitted, and Diaf et al. [73] apply a comparable techno-economic framework to a hybrid PV/wind system on Corsica Island, evaluating designs through the levelised cost of electricity (LCOE). While simple, this matching approach often produces oversized systems with unknown economic returns, particularly when feed-in restrictions prevent the sale of surplus energy. An alternative approach minimises the energy lost to feed-in restriction, but this can lead to undersized systems that do not significantly reduce public grid reliance [7]. Both methods lack a structured optimisation framework that balances system size, economic return, and sustainability objectives.

Various optimisation methods have been presented for PV and wind energy systems, applying diverse objective functions grouped into engineering and economic categories. Ning et al. [74] proposed a Building Information Modelling (BIM)-based automatic design method minimising the cost-to-power ratio. Liu et al. [71] analysed PV energy systems with a simulation and optimisation method that accounts for carbon dioxide emissions and economic metrics. Diaf et al. [73] presented a technical-economic optimisation using the LCOE. Fan and Xia [72] proposed a multi-objective optimisation model for retrofitting buildings with PV energy systems using NPV and payback period. Energy storage integration has been addressed by Mehrabankhomartash et al. [75], who analysed storage with customer damage functions, and Starke et al. [76], who proposed a methodology for concentrated solar power plants with thermal storage. Commercially available tools such as HOMER [77] and PVSyst [78] have also been used for system design, along with geographical information system (GIS)-based workflows [79].

A specific sub-category of energy system design concerns semi-autonomous structures, buildings that aim to minimise their reliance on the public energy grid. The smart-home review of Lashkari et al. [80] provides the broader taxonomy of energy management strategies under which such designs sit, and the off-grid case study of Akinyele et al. [81] illustrates the operational regime that semi-autonomous structures approximate as grid connectivity is reduced. Waterside semi-autonomous structures, such as floating homes, present a particularly challenging case: they must rely on the grid as little as possible because maintaining a stable grid connection at waterside locations is difficult and costly. Gorzka et al. [82] reported a case study of zero-emission, off-grid autonomous houseboats at selected European locations in *Polish Maritime Research*, examining the feasibility of cutting the houseboat free from the public grid altogether. Grela et al. [83] subsequently described an e-houseboat that integrates RES with a hybrid energy storage system (ESS) to enhance energy autonomy. Compared to regular semi-autonomous structures on land, floating homes introduce additional constraints: the structure can be rotated within its mooring and design constraints (enabling PV orientation optimisation), the insulation and sun exposure differ (affecting thermal modelling and air-conditioning system loads), and the physical structure has limited weight capacity (constraining battery sizing). Ishaque et al. [84] proposed a structural design for a low-cost floating-home case study with safety as the primary criterion, illustrating the load-bearing and rotational considerations that arise when a habitable structure must remain stable on water.

Two gaps persist in the literature on physical PV energy system design. First, no consumption-aware optimisation method exists for multi-string PV energy systems under feed-in restrictions that offers per-string convergence guarantees for any objective and tariff whose unimodality in string power can be verified. Second, no dedicated methodology addresses the specific constraints of waterside semi-autonomous structures, including coupled thermal-electrical modelling, stochastic load profile estimation, and joint PV-storage optimisation under weight limitations.

1.5 Research objectives and contributions

The three interconnected challenges identified above are addressed through three theses, each comprising two or three points.

Objective 1. *To develop machine learning methods for accurate day-ahead EPF.*

Accurate short-term EPF is a prerequisite for the profitable operation of battery-enhanced renewable power plants. This objective encompasses:(a) the development of a novel general-purpose stopping criterion that improves NN generalisation by jointly monitoring training and validation loss dynamics;(b) the systematic design of LSTM networks for day-ahead EPF, combining grid search with hyperparameter optimisation and integrating the Correlation-Driven Stopping Criterion (CDSC) into the optimisation pipeline; and(c) the comparative evaluation of LSTM and Transformer-based foundation model architectures for direct multi-step forecasting across multiple markets.

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.*

The economic viability of battery-enhanced renewable power plants depends on the trading strategy employed on the DAM. This objective involves: (a) the development of a parametrised mathematical model of DC-coupled hybrid inverter systems and a novel adaptive gradient-descent-based energy trading optimisation method; and (b) the development of a differential-evolution-based metaheuristic for the same domain.

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.*

Industrial and household-size prosumers face diverse constraints that render conventional PV and ESS sizing methods suboptimal. This objective encompasses: (a) the development of a consumption-aware dichotomous search-based string optimisation algorithm for multi-string PV energy systems under feed-in restrictions; and (b) the development of a comprehensive energy system design methodology for semi-autonomous structures.

Taken together, the three objectives form a coherent decision-support framework for battery-enhanced renewable power plants, spanning *prediction*, *control*, and *design*. Thesis I predicts day-ahead prices; Thesis II converts those predictions into an optimised trading strategy; and Thesis III designs and sizes the physical battery-enhanced PV energy system on which that strategy operates. The three stages are coupled: accurate forecasts raise the profit extracted from a given system, while the physical design sets the envelope of what any trading strategy can achieve. Figure 1.1 presents this prediction–control–design chain as a graphical abstract of the dissertation.

1.6 Structure of the dissertation

The remainder of this dissertation is organised as follows:

Chapter 2 addresses Thesis I and presents the machine learning methods developed for day-ahead EPF. The chapter covers the CDSC as a general-purpose training regularisation method (Thesis I.1), the design of LSTM networks for EPF, combining grid search with hyperparameter optimisation (Thesis I.2), and the evaluation of Transformer-based foundation models with optimised context length for direct multi-step forecasting (Thesis I.3).

Chapter 3 addresses Thesis II and presents the energy trading strategy optimisation methods for battery-enhanced renewable power plants. The chapter introduces a parametrised system model for DC-coupled hybrid inverter systems and the adaptive gradient-descent-based OptEnT-GD algorithm (Thesis II.1), followed by the differential-evolution-based OptEnT-DE algorithm and a comparative evaluation that leads to a system sizing guide (Thesis II.2).

Chapter 4 addresses Thesis III and presents the optimisation and modelling methods for the physical sizing of PV energy systems. The chapter covers the consumption-aware dichotomous search-based string optimisation of multi-string PV energy systems (Thesis III.1) and the energy system design methodology for waterside semi-autonomous structures (Thesis III.2).

Chapter 5 consolidates the formal thesis statements for all seven sub-theses (I.1–III.2) together with the source publications that anchor each one.

Finally, Chapter 6 summarises the main findings, revisits the thesis points with their supporting evidence, and outlines directions for future research.

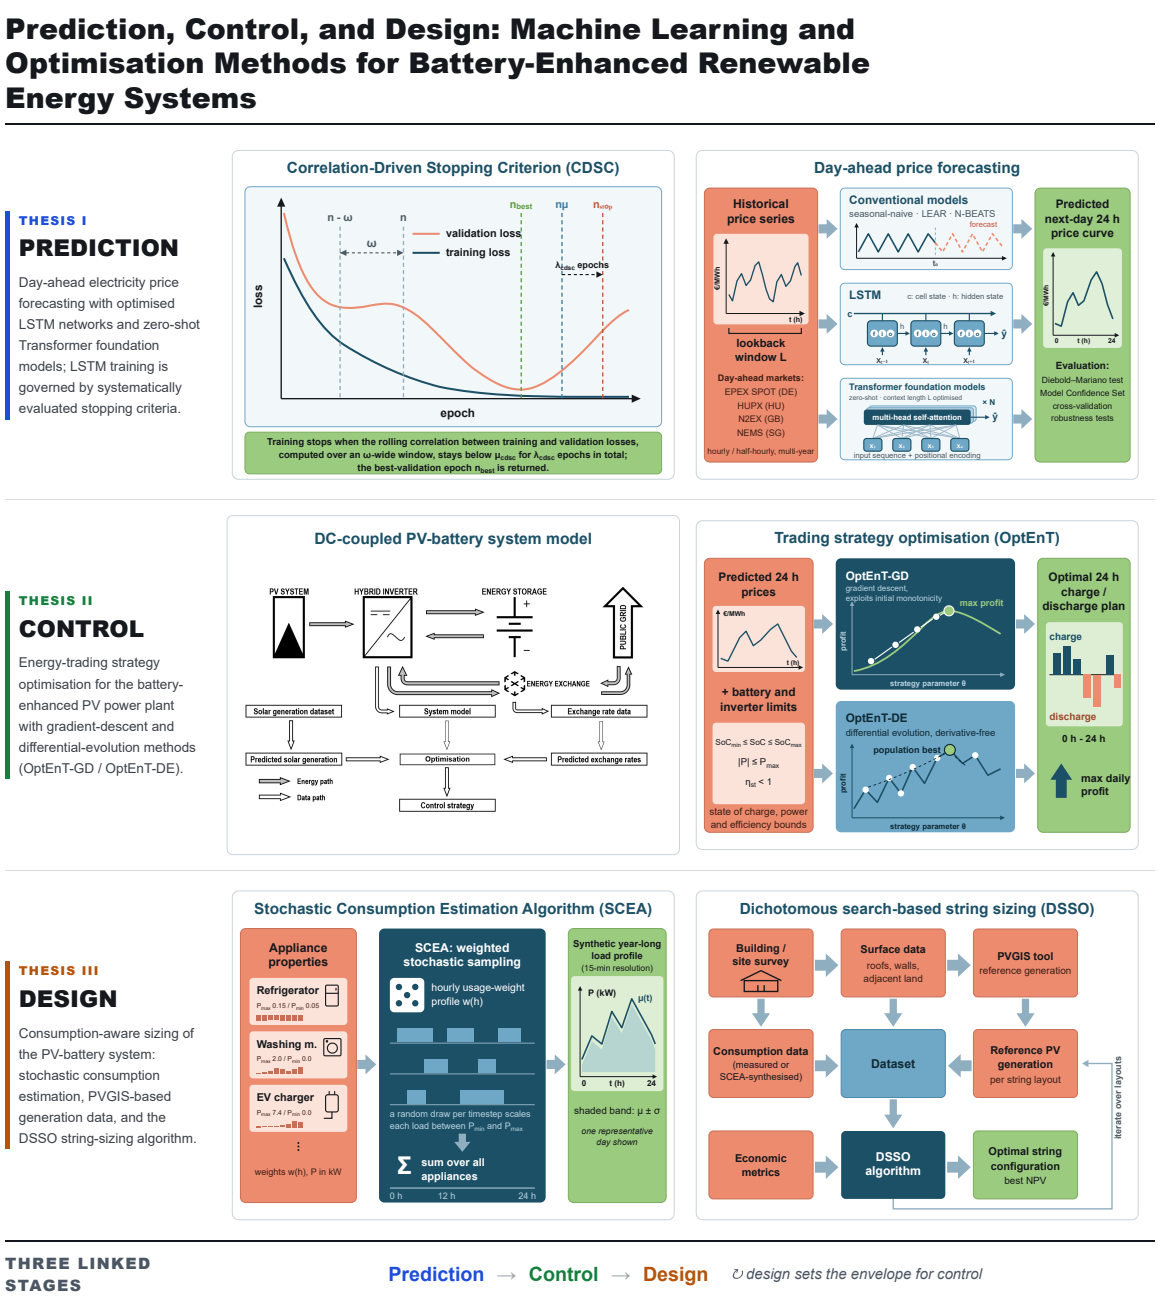


Figure 1.1: Structure description of the dissertation: the prediction–control–design decision-support pipeline for battery-enhanced renewable power plants. Thesis I forecasts day-ahead prices, Thesis II optimises the trading strategy from those forecasts, and Thesis III designs the physical battery-enhanced PV energy system on which the strategy operates.

Chapter 2

Day-Ahead Electricity Price Forecasting Using Machine Learning

2.1 Introduction

As established in Section 1.2, accurate day-ahead electricity price forecasting (EPF) is essential for the profitable operation of battery-enhanced renewable power plants and for rational decision-making in competitive electricity markets. This chapter presents the machine learning methods developed for this task, spanning three interconnected contributions: I proposed a novel training stopping criterion applicable to neural networks (NNs) in general (Thesis I.1), I performed a systematic design and hyperparameter optimisation of Long Short-Term Memory (LSTM) networks for price forecasting on multiple day-ahead markets (DAMs) (Thesis I.2), and I evaluated pre-trained Transformer-based foundation models with systematic context-length optimisation (Thesis I.3).

The stopping criterion determines when the training process terminates and is a hyperparameter whose influence on generalisation is well recognised in the deep-learning literature. Prechelt catalogued and empirically compared several patience-based variants of early stopping on synthetic and real tasks, framing the question of *when* to halt training as a practical concern with consequences for both training time and final model quality [31]. Goodfellow et al. subsequently presented early stopping as the canonical regularisation recommendation in their textbook treatment of deep learning, characterising it as one of the most commonly used forms of regularisation and relating its effect to that of an implicit weight-decay penalty [85]. The simplest approach, training for a fixed maximum number of epochs, in its naive form provides no mechanism for model selection based on out-of-sample performance, frequently resulting in overfitted models. Classical early stopping monitors the validation loss and halts training when the validation loss fails to improve for a specified patience period [31]. While widely adopted and effective as a baseline regulariser, early stopping monitors only the validation loss curve and is sensitive to its patience hyperparameter, which is rarely tuned systematically.

Several extensions have been proposed to address the limitations of classical early stopping. One line of work has refined the stopping decision for individual models. Iyer and Rhinehart presented a statistical method that triggers once the moments of the training error signal become stationary [86]. Wu and Liu introduced a fuzzy-based early stopping algorithm aimed at improving generalisation [87]. Lalis et al. developed an adaptive stopping criterion for backpropagation training of feed-forward networks [88]. Other studies have examined ensemble settings. Ennett, Frize, and Scales evaluated the logarithmic-sensitivity index as a stopping criterion for neural networks predicting rare outcomes [89]. Nguyen, Abbass, and McKay studied stopping criteria for ensembles of evolutionary artificial neural networks, including the minimum validation fitness of the ensemble [90]. Adjacent domains have also contributed. Zhu et al. proposed confidence-based stopping criteria for active learning over data annotation [91]. Altschuler and Bloodgood used the predicted change of the F-measure to stop active learning for text classification [92]. Ishibashi and Hino derived a stopping criterion for active learning from deterministic generalisation bounds rooted in Probably Approximately Correct (PAC)-Bayesian theory [93]. Finally, gradient-based criteria have been proposed to enable early stopping without a validation set altogether [94].

Despite these contributions, most published stopping criteria are not implemented in

popular machine learning frameworks, limiting their adoption. The maximum number of epochs and classical early stopping remain the dominant strategies in practice [31]. In contrast, hyperparameter tuning of the stopping criterion receives far less attention than that of the model architecture [8].

The application of LSTM networks [25] to day-ahead EPF has been explored in several studies. Ugurlu et al. [26] applied recurrent neural networks (RNNs), including LSTMs, to EPF and demonstrated their superiority over feed-forward architectures. Kuo and Huang [27] proposed a hybrid structured deep NN for EPF. Peng et al. [95] applied a long short-term memory model trained with a differential evolution (DE) algorithm to electricity price prediction. The design choices for LSTM-based EPF models, including lookback window length, network depth and width, activation function, normalisation strategy, and regularisation technique, are typically investigated in isolation [26] rather than jointly optimised. Furthermore, the interaction between the stopping criterion and these architectural choices has not been systematically explored: it remains unclear whether the optimal architecture differs when a correlation-based stopping criterion is used instead of standard early stopping.

The emergence of pre-trained time-series foundation models represents a paradigm shift in forecasting [37]. Unlike task-specific models that require training on each new dataset, such foundation models are pre-trained on large heterogeneous corpora and can generate forecasts in a zero-shot regime. Chronos tokenises continuous time-series values, after per-series scaling, into a fixed integer vocabulary and trains an encoder–decoder transformer with a language-model-style cross-entropy objective over the resulting token distribution [37]. Chronos-Bolt retains the encoder–decoder backbone but replaces the per-token autoregressive head with a patch-based decoder that directly emits multi-step quantile forecasts, reported to deliver order-of-magnitude faster inference at comparable forecasting accuracy [38]. TimesFM exemplifies the same paradigm with a patched decoder-only architecture [39]. Moirai is presented as a unified model for universal time-series forecasting and employs a masked-encoder design [40]. The context window length, the number of historical observations provided as input, is the primary hyperparameter governing their performance. Hornek et al. [41] conducted an early benchmark of foundation models on EPF tasks, finding that they perform on par with traditional approaches but do not consistently outperform well-tuned task-specific models. However, the benchmarked models have since been updated to more modern versions, and their sensitivity to context length has not been systematically optimised on day-ahead electricity market data across markets with different temporal resolutions and volatility profiles.

The choice of cross-validation (CV) strategy during hyperparameter optimisation directly affects the quality of the selected model configuration [43]. For time-series data, standard k-fold CV is inappropriate because it violates the temporal ordering assumption [43]. Temporally sound alternatives that preserve the chronological order of observations include holdout, expanding-window, and sliding-window schemes. Tashman's review of out-of-sample tests of forecasting accuracy laid out the rolling-origin evaluation protocol that replaces randomly assigned folds with chronologically ordered forecast origins, and surveyed the methodological choices that arise in setting up such tests [42]. Cerqueira et al. subsequently presented an empirical study of performance-estimation methods for time-series forecasting models, including CV variants and prequential schemes [44]. When lag-based features are used, purged and embargoed CV, which removes observations near fold boundaries with an embargo width of $L + n$ (lookback window plus prediction horizon), prevents look-ahead contamination [45]. Whether a single-fold extended-validation scheme suffices or a multi-fold leave-one-year-out design is required for robust model selection in EPF has not been investigated within a unified hyperparameter optimisation framework.

This chapter addresses these gaps. Section 2.2 presents the novel Correlation-Driven Stopping Criterion (CDSC) as a general-purpose method and evaluates it across diverse machine learning tasks beyond electricity price forecasting. Section 2.3 describes the joint optimisation of LSTM networks over architecture, lookback window length, and stopping-criterion parameters using the Tree-structured Parzen Estimator (TPE) [96] at high-performance computing (HPC) scale, with systematic comparison of CV strategies. Section 2.4 evaluates foundation Transformer models with exhaustive context-length optimisation and benchmarks them against the best LSTM configurations. The formal thesis statements for Thesis I are gathered in Section 5.1, alongside those of Thesis II and Thesis III.

2.2 Correlation-Driven Stopping Criterion for neural network training

2.2.1 Motivation and context

Classical early stopping and the maximum number of epochs criterion remain the dominant stopping strategies in NN training [31], despite a range of proposed alternatives. Iyer and Rhinehart introduced a statistical stopping criterion that triggers once the statistical moments of the training error signal become stationary [86]. Wu and Liu proposed a fuzzy-based early stopping algorithm aimed at improving NN generalisation [87]. Lalis et al. described an adaptive stopping criterion for backpropagation training of feed-forward NNs [88]. Mahsereci et al. introduced a stopping criterion based on local statistics of the computed gradients, removing the reliance on a held-out validation set [94]. A key limitation shared by all validation-loss-based methods is that they monitor only a single metric: the validation loss. With adequate parameters, such methods can recognise when the validation loss stops improving, but they cannot determine whether the optimisation algorithm has reached a temporary plateau or whether the model's generalisation performance can still improve [31]. The change in the training loss is ignored entirely.

This limitation motivates a different approach. During well-conditioned training, the training and validation losses decrease together; when the model begins to overfit, the training loss continues to decrease while the validation loss stagnates or increases [85]. The divergence between the two loss curves is therefore a direct indicator of overfitting. By monitoring the joint dynamics of both curves, rather than the validation loss alone, a richer and earlier signal for overfitting detection becomes available.

The Correlation-Driven Stopping Criterion, proposed as part of this dissertation, operationalises this idea by tracking the rolling Pearson correlation coefficient between the training and validation loss curves. Training is terminated when this correlation falls below a predefined threshold, indicating that the two curves have begun to diverge. Unlike the methods reviewed in Section 2.1, the CDSC is designed as a general-purpose regularisation method, not specific to any particular application domain. To demonstrate this generality, it was evaluated across four benchmark problems (binary classification, multiclass classification, multivariate regression, and univariate time-series prediction) using three NN architectures (feed-forward, convolutional, and LSTM), and compared against early stopping and the maximum number of epochs criterion. Its application to electricity price forecasting is addressed separately in Thesis I.2.

2.2.2 Method

The key mechanic of the newly proposed CDSC method is that it monitors the rolling Pearson correlation coefficient between the training and validation loss curves within a sliding window of size ω. The rolling Pearson correlation coefficient from epoch $n_{i-\omega+1}$ to epoch n_i, an ω-epoch window, is calculated as:

$$r = \frac{\sum_{j=n_{i-\omega+1}}^{n_i} \left(e_{\text{tr}}^{(j)} - \overline{e}_{\text{tr}}\right)\left(e_{\text{va}}^{(j)} - \overline{e}_{\text{va}}\right)}{\sqrt{\sum_{j=n_{i-\omega+1}}^{n_i} \left(e_{\text{tr}}^{(j)} - \overline{e}_{\text{tr}}\right)^2 \sum_{j=n_{i-\omega+1}}^{n_i} \left(e_{\text{va}}^{(j)} - \overline{e}_{\text{va}}\right)^2}}, \tag{2.1}$$

where n_i is the i-th epoch ($i > \omega$, $\omega \geq 1$), $e_{\text{tr}}^{(j)}$ and $e_{\text{va}}^{(j)}$ are the training and validation errors at the j-th epoch, and $\overline{e}_{\text{tr}}$ and $\overline{e}_{\text{va}}$ denote the mean training and validation errors within the window, respectively. The monitored loss can be chosen arbitrarily; any numeric loss function on which the Pearson correlation can be computed is suitable.

When the problem is well-conditioned (i.e., the dataset is adequate, the model fits the task, and the hyperparameters are reasonable), the rolling correlation starts high and decreases over the course of the epochs as the model begins to overfit. Training is terminated when the rolling correlation falls below a predefined threshold μ_{cdsc}, and the model with the best validation error up to that point is selected. The epoch at which the correlation first falls below μ_{cdsc} is denoted n_μ.

A patience parameter λ_{cdsc} relaxes this strict condition: training continues until the rolling correlation has been below μ_{cdsc} for λ_{cdsc} (not necessarily consecutive) epochs. The epoch at which training stops is denoted n_{stop}, and the epoch with the best validation error is n_{best}.

The key distinction between the CDSC and early stopping lies in the information each method uses. Early stopping monitors only the validation loss; with adequate parameters, it can recognise when the validation loss stops improving, but it cannot determine whether the optimisation algorithm has reached a plateau or whether further improvement is possible. The CDSC monitors both training and validation errors to assess the state of the training process. If the two metrics move in sync, the model's out-of-sample performance can improve further. If they start to diverge, the model is learning training-set-specific features, which is a sign of overfitting. This dual-monitoring approach also provides diagnostic value: noisy correlation curves and low initial correlation values indicate that the problem formulation, the model hyperparameters, or the training configuration should be reconsidered.

Figure 2.1 illustrates the CDSC. The rolling Pearson correlation with window size ω is computed stepwise during training. At epoch n_μ, the correlation coefficient falls first below the threshold μ_{cdsc}, but according to the relaxed stopping criterion, the training process continues. At epoch n_{stop}, the training stops as the evaluated rolling correlation coefficient between the training loss and the validation loss was below the correlation threshold μ_{cdsc} in λ_{cdsc} epochs. Then, the model with the best validation error is selected, marked by n_{best}.

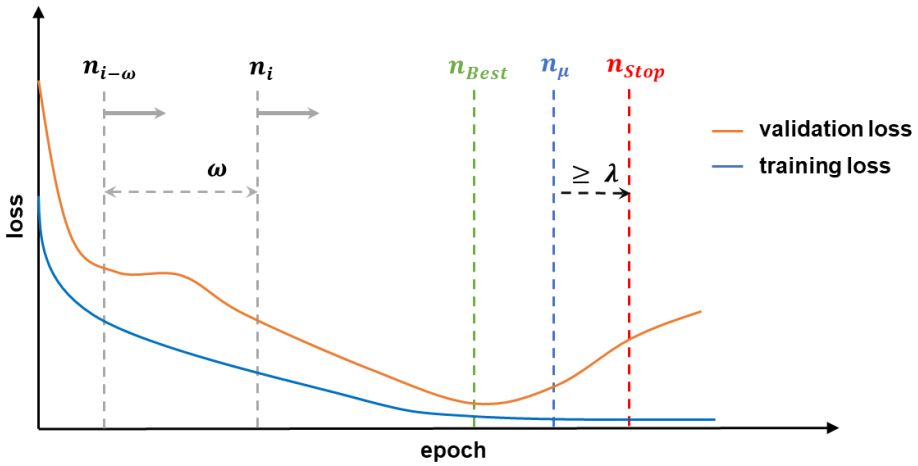


Figure 2.1: Visual representation of the Correlation-Driven Stopping Criterion. The rolling Pearson correlation coefficient of the training and the validation error is calculated for ω-wide windows step-by-step. n_μ denotes the epoch where the rolling correlation coefficient falls first below the threshold μ_{cdsc}. The training stops at epoch n_{stop}, and n_{best} denotes the epoch of the best validation error selected.

The CDSC algorithm is summarised in Algorithm 1. Its computational complexity is $O(\omega)$ per epoch, which does not significantly increase the total training time.

In the pseudocode, the input parameter λ_{cdsc} initialises a separate runtime counter c, which is decremented each time the rolling correlation falls below μ_{cdsc}; training stops once c reaches zero or below, so that $\lambda_{\text{cdsc}} = 0$ stops at the first sub-threshold epoch. The algorithm also receives the maximum number of epochs N as an input parameter. This is only necessary so that the learning process does not run indefinitely; the value of N does not need to be restricted significantly, and it is sufficient to choose a sufficiently large value.

Algorithm 1: CDSC: Correlation-Driven Stopping Criterion. Training stops once the rolling correlation between the training and validation losses has been below the threshold μ_{cdsc} for a cumulative total of λ_{cdsc} epochs, and the epoch with the best validation error (n_{best}) is returned.

Input: ω window size
μ_{cdsc} correlation threshold
λ_{cdsc} patience
N maximum number of training epochs

Output: n_{best} best epoch
n_{stop} stop epoch

Initialise $n_{\text{best}} = \texttt{NULL}$, $n_{\text{stop}} = \texttt{NULL}$;
Initialise counter $c = \lambda_{\text{cdsc}}$;
Initialise training and validation error lists: $e_\text{tr} = []$, $e_\text{va} = []$;
for $i = 1$ **to** N **do**
    Retrieve $e_{\text{tr}}^{(i)}$ and $e_{\text{va}}^{(i)}$;
    Append to e_tr and e_va;
    if $i > \omega$ **then**
        $r \leftarrow \text{corr}\big(e_\text{tr}[i{-}\omega : i],\ e_\text{va}[i{-}\omega : i]\big)$;
        if $r < \mu_{\text{cdsc}}$ **then**
            $c \leftarrow c - 1$;
            if $c \le 0$ **then**
                $n_{\text{stop}} \leftarrow i$;
                $n_{\text{best}} \leftarrow \arg\min_{k \le i} e_\text{va}[k]$;
                stop training;

`// Criterion never triggered: fall back to the epoch budget`
if $n_{\text{stop}} = \texttt{NULL}$ **then**
    $n_{\text{stop}} \leftarrow N$;
    $n_{\text{best}} \leftarrow \arg\min_{k \le N} e_\text{va}[k]$;
return $n_{\text{best}}, n_{\text{stop}}$

2.2.3 Experimental methodology

The CDSC was evaluated as a general-purpose method across four benchmark problems selected for diversity in task type, class balance, dataset size, and problem domain:

1. **Credit Card Fraud Detection** [97]: binary classification, highly unbalanced (0.172% positive), 284 807 samples, 30 features. Model: feed-forward neural network (FFNN) with two hidden layers of 256 neurons, trained with Adam optimiser, batch size 512, 250 epochs.

2. **MNIST Handwritten Digit Classification** [98]: multiclass classification (10 classes), balanced, 70 000 samples, 784 features. Model: convolutional neural network (CNN) with two 32-unit convolutional layers (3×3 kernel, 2×2 max pooling), batch size 256, 150 epochs.

3. **Boston Housing Price Regression**: multivariate regression, 506 samples, 13 features. Model: FFNN with two hidden layers of 256 neurons, batch size 8, 250 epochs.

4. **Gold Daily Price Change Prediction**: univariate time-series regression, 9823 samples, 14 lagged features. Model: LSTM with 64 units, batch size 64, 150 epochs.

The NN models used for each benchmark problem are summarised in Table 2.1.

Table 2.1: NN architectures used for the CDSC benchmark evaluation. The four models cover the FFNN, CNN, and LSTM architectures, exercising the criterion across binary classification, multiclass classification, regression, and time-series tasks rather than a single architecture.

Dataset	Model type	Input layer	Hidden layer(s)	Output layer
Credit Card Fraud Detection	FFNN	30	2×256 neurons ReLU	1 neuron sigmoid
MNIST	CNN	784	2×32 conv. units 3×3 kernel 2×2 max pooling	10 neurons softmax
Boston Housing	FFNN	13	2×256 neurons ReLU	1 neuron linear
Gold Daily Price Change	LSTM	14	1×64 LSTM units linear	1 neuron linear

A rigorous evaluation methodology was applied to ensure fair comparison. For each dataset, 50 independently shuffled train/validation/test splits (80%/10%/10%) were created. All models were trained for a sufficiently large number of epochs, N, producing complete epoch-level logs recording training, validation, and test errors at each epoch. The three stopping criteria were then simulated on these shared logs by truncating at the epoch specified by each criterion, eliminating confounding effects from different random initialisations. This approach ensures that all stopping criteria are evaluated under identical conditions, with the same initialisations and learning processes, and that the comparison reflects only the quality of the stopping decision itself.

Figure 2.2 illustrates the evaluation methodology. The models were always trained to the full N, and the stopping criteria were applied post hoc to the saved epoch logs. This eliminates the need to retrain models for each stopping-criterion configuration and enables exhaustive exploration of the parameter space.

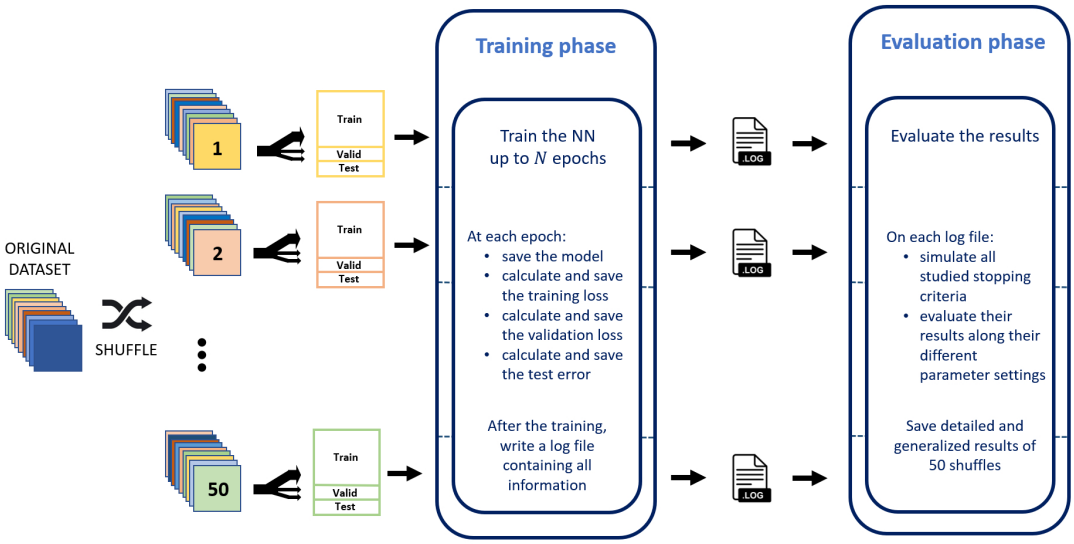


Figure 2.2: Graphical representation of the comparative evaluation methodology. For each of the 50 shuffled datasets, models are trained for N epochs, producing complete epoch logs. The three stopping criteria are simulated on these shared logs by truncation, ensuring identical training conditions across all comparisons.

The following evaluation metrics were used to compare the stopping criteria:

- *Mean percentage error (*MPE*) change* ($\overline{\%e}_{\text{ts}}^{S(\text{pset})}$): the mean relative change of the test error achieved with stopping criterion S and parameter set pset compared to the baseline (model selected at the best validation error across the full N epochs).
- *Mean stopping epoch* ($\overline{n_{\text{stop}}}^{S(\text{pset})}$): the average epoch at which training was terminated.
- *Mean selected epoch* ($\overline{n_{\text{best}}}^{S(\text{pset})}$): the average epoch at which the best validation error was observed.

The CDSC parameter space was explored via grid search: $\omega \in \{5, 10, \ldots, 50\}$, $\lambda_{\text{cdsc}} \in \{0, 5, \ldots, N-50\}$, and $\mu_{\text{cdsc}} \in \{-1.00, -0.99, \ldots, 1.00\}$. The grid search parameters are summarised in Table 2.2. The primary evaluation metric was the mean percentage error change relative to the baseline model (the best-validation checkpoint of the full N-epoch run, as defined above).

Table 2.2: Parameter space for the CDSC grid search. Window size and patience are scanned in steps of 5, whereas the correlation threshold is swept across its full range $[-1, 1]$ at a finer resolution (step 0.01), the parameter to which the stopping decision proves most sensitive.

Parameter	Lower bound	Upper bound	Step size
Window size (ω)	5	50	5
Patience (λ_{cdsc})	0	$N - 50$	5
Threshold (μ_{cdsc})	-1	1	0.01

The importance of the CDSC parameters (ω, λ_{cdsc}, μ_{cdsc}) was determined by fitting hyperparameter-tuned random forest (RF) classification models on the results of the grid search trials. The classification task was formulated as a binary problem: grid search trials that improved the baseline model performance were labelled as "good", and all others as "bad". The input features were the values of the ω, λ_{cdsc}, and μ_{cdsc} parameters.

An automated tuning procedure was also evaluated to reduce the cost of the grid search. Bayesian optimisation (BO) with gradient-boosted regression trees (GBRT) as the surrogate function was applied, following the methodology of Makarova et al. [99] and Ishibashi et al. [100]. In the analysis, 50 initial points were calculated before the first approximation of the surrogate function, and the negative expected improvement was used for minimisation. The tuning process was stopped when the latest parameter setting yielded a regret score comparable to that of one of the best 50 parameter settings. The process was repeated 50 times per problem to mitigate the stochasticity of both NN training and hyperparameter tuning.

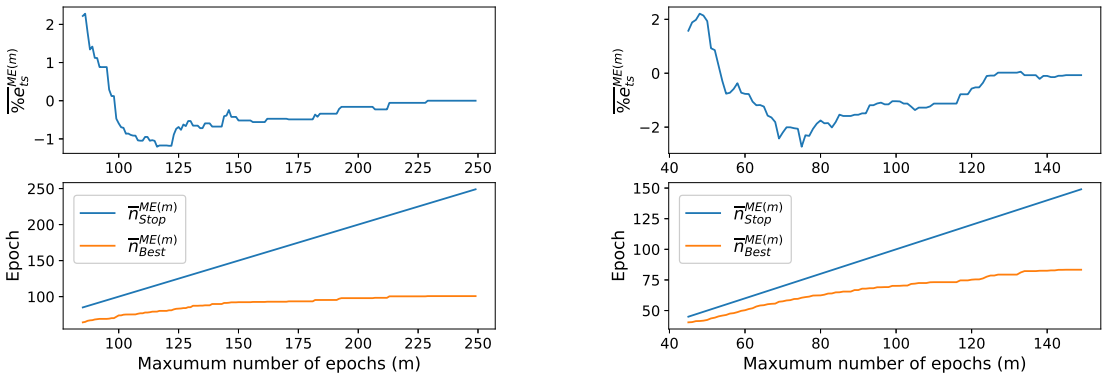
To compare the three stopping criteria in a more general fashion and to obtain information on their overall stability, a practical comparative regret-bound analysis was performed. The regret measures the "cost" of using a stopping criterion in real-world conditions compared to the best-case scenario, where the test error is known. The regret for the CDSC was calculated as:

$$\text{regret} = \frac{1}{50} \sum_{k=1}^{50} \frac{e_{\text{ts}_k}^{CD(\omega,\lambda_{\text{cdsc}},\mu_{\text{cdsc}})} - e_{\text{ts}_k}^{*}}{e_{\text{ts}_k}^{*}} \times 100, \tag{2.2}$$

where $e_{\text{ts}_k}^{CD(\omega,\lambda_{\text{cdsc}},\mu_{\text{cdsc}})}$ is the test error achieved with the CDSC for the k-th shuffled dataset with the given parameter setting, and $e_{\text{ts}_k}^{*}$ is the minimum test error observed during the entire training process. Analogous formulations were used for the maximum number of epochs and early stopping criteria.

2.2.4 Results

The maximum number of epochs stopping criterion (benchmarked here as an epoch budget with best-validation checkpoint selection) stops training when a fixed epoch m ($m = 1, \ldots, N$) is reached, and the model with the best validation error up to epoch m is selected. Figure 2.3 shows the effect of the parameter m on the performance of this criterion across the four benchmark datasets. The top panels in each subfigure show the mean percentage error change as a function of m, while the bottom panels show the average stopping and selected epochs.



(a) Credit Card Fraud Detection

(b) MNIST

(c) Boston Housing

(d) Gold Daily Price Change

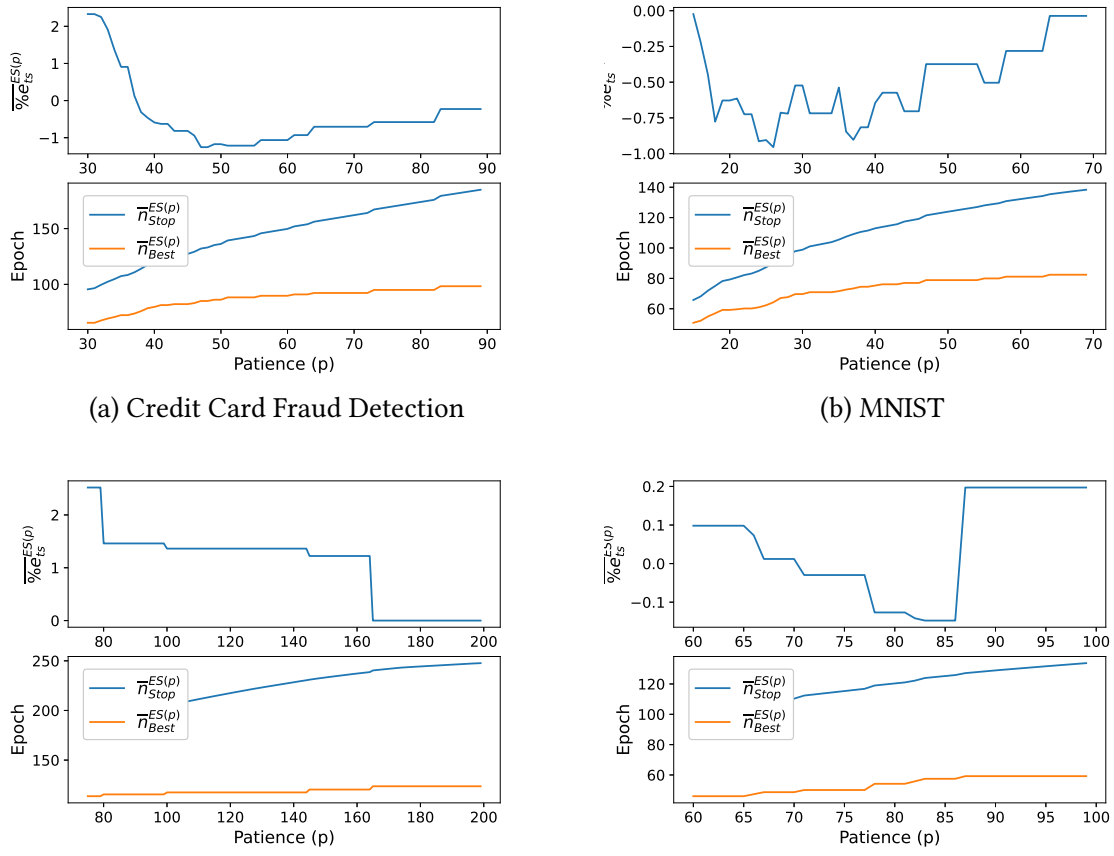
Figure 2.3: Effect of the maximum epoch parameter m on the performance of the maximum number of epochs stopping criterion. Top panels: mean percentage error change ($\overline{\%e_{\text{ts}}}^{ME(m)}$) as a function of m. Bottom panels: average stopping epoch ($\overline{n_{\text{stop}}}^{ME(m)}$) and average selected epoch ($\overline{n_{\text{best}}}^{ME(m)}$). An intermediate m lowers the test error below the full-budget baseline on all four datasets, confirming that the epoch budget itself acts as a regularising hyperparameter.

The results with the best parameter m_{Best} for each dataset are summarised in Table 2.3. The criterion improved out-of-sample performance on all four datasets, with the greatest improvement on MNIST (−2.73%).

Table 2.3: Metrics for the maximum number of epochs criterion with the best parameter m_{Best}. Capping the epoch budget lowered the test error on all four datasets, most on MNIST (−2.73%) and only marginally on Boston Housing and Gold.

Dataset	m_{Best}	$\overline{n_{\text{stop}}}$	$\overline{n_{\text{best}}}$	$\overline{\%e_{\text{ts}}}^{ME(m_{\text{Best}})}$
Credit Card Fraud Det.	116	116.00	79.22	−1.20
MNIST	75	75.00	60.48	−2.73
Boston Housing	186	186.00	113.64	−0.18
Gold Daily Price Change	103	103.00	45.66	−0.05

Early stopping stops training when the validation error does not improve for p_{es} consecutive epochs. Figure 2.4 shows the effect of the patience parameter p_{es} on performance. On the MNIST and Credit Card Fraud Detection problems, the mean percentage error change was negative over a large portion of the parameter space. On the Boston Housing Price Regression problem, the mean percentage error change was positive or zero across all tested patience values.



(a) Credit Card Fraud Detection

(b) MNIST

(c) Boston Housing

(d) Gold Daily Price Change

Figure 2.4: Effect of the patience parameter p_{es} on the performance of the early stopping criterion. Top panels: mean percentage error change ($\overline{\%e_{\text{ts}}}^{ES(p_{\text{es}})}$) as a function of p_{es}. Bottom panels: average stopping and selected epochs. Early stopping reduced the test error over a wide patience range on MNIST and Credit Card Fraud Detection, but no patience value improved on the baseline for Boston Housing, where the validation loss plateaus.

The best early stopping results are summarised in Table 2.4. The most significant improvement was observed on the Credit Card Fraud Detection dataset (−1.25%).

Table 2.4: Metrics for the early stopping criterion with the best patience parameter $p_{\text{es,Best}}$. The largest gain was on Credit Card Fraud Detection (−1.25%), while Boston Housing saw no improvement at any patience.

Dataset	$p_{\text{es,Best}}$	$\overline{n_{\text{stop}}}$	$\overline{n_{\text{best}}}$	$\overline{\%e}_{\text{ts}}^{ES(p_{\text{es,Best}})}$
Credit Card Fraud Det.	47	132.10	85.10	−1.25
MNIST	26	90.24	64.24	−0.96
Boston Housing	165	240.30	123.64	0.00
Gold Daily Price Change	83	123.94	57.48	−0.15

Since the CDSC has three parameters, a grid search was used to explore the parameter space comprehensively (see Table 2.2). The best parameter combinations for each dataset are shown in Table 2.5. The CDSC improved performance on every tested problem, with the mean percentage error change negative in every case. The largest improvement was observed on MNIST (−2.89%). These best settings are identified retrospectively on the test error, so the reported figures are the gain achievable under oracle parameter selection rather than an out-of-sample estimate; the BO study below addresses how reachable such settings are in practice.

Table 2.5: Metrics for the CDSC with the best parameter combination ($\mu_{\text{cdsc,Best}}$, ω_{Best}, $\lambda_{\text{cdsc,Best}}$) for each dataset. The CDSC lowered the test error on every dataset, most on MNIST (−2.89%); the best threshold is positive for the cleaner Credit Card Fraud Detection and MNIST problems and negative for the smaller, noisier Boston Housing and Gold problems.

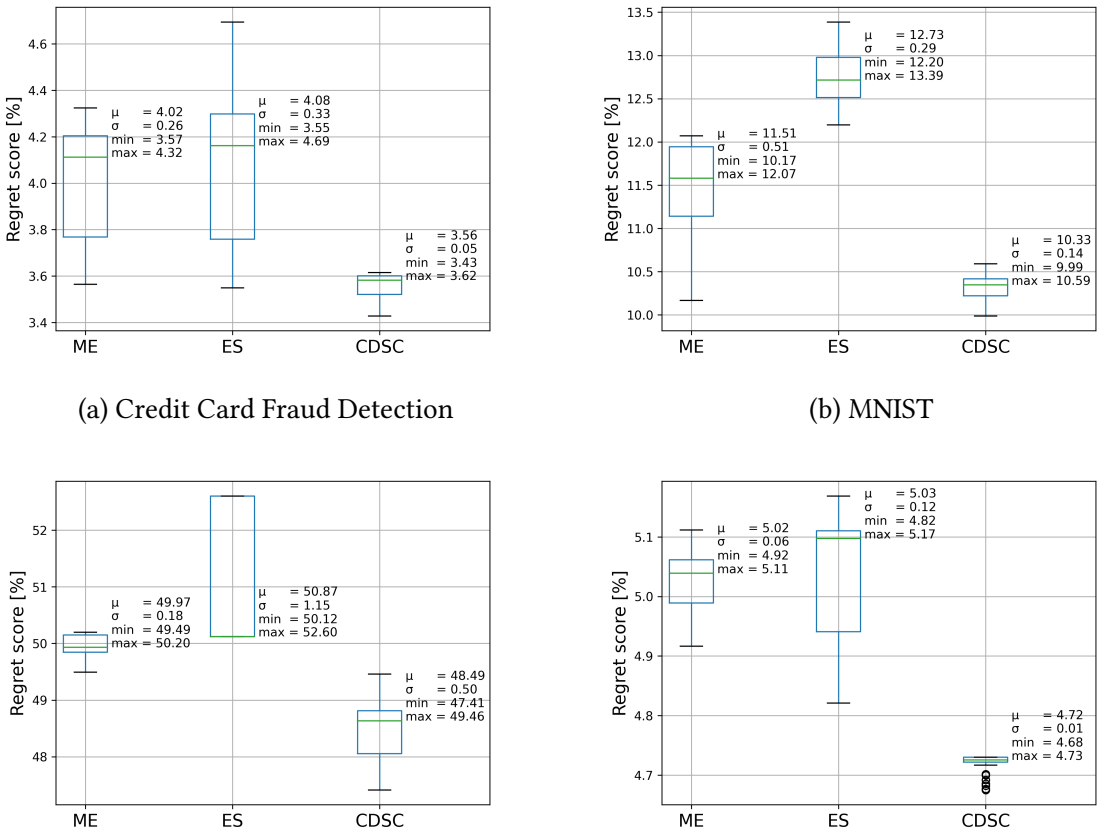
Dataset	Best parameters			$\overline{n_{\text{stop}}}$	$\overline{n_{\text{best}}}$	$\overline{\%e}_{\text{ts}}$
	μ_{cdsc}	ω	λ_{cdsc}			
Credit Card Fraud Det.	0.75	30	65	119.28	79.00	−1.35
MNIST	0.40	35	10	76.30	61.26	−2.89
Boston Housing	−0.40	10	25	197.76	108.28	−1.56
Gold Daily Price Change	−0.15	30	25	99.54	39.22	−0.27

The comparative results across the three stopping criteria are summarised in Table 2.6. The CDSC with optimal parameters outperformed both baseline methods on all four datasets. The improvement was most pronounced on the Boston Housing dataset, where the CDSC achieved a 1.56% improvement in out-of-sample performance while early stopping yielded no improvement and the maximum number of epochs criterion achieved only 0.18%.

The regret distributions for the 50 best-performing parameter settings of each criterion are shown in Figure 2.5. In all four benchmark problems, the mean, minimum, and standard deviation of the regret achieved with the CDSC were lower than the corresponding statistics for the other two criteria.

Table 2.6: Mean percentage error change relative to baseline for the three stopping criteria with their best parameter settings. Negative values indicate improved out-of-sample performance.

Dataset	Max. epochs	Early stopping	CDSC
Credit Card Fraud Detection	−1.20	−1.25	**−1.35**
MNIST	−2.73	−0.96	**−2.89**
Boston Housing	−0.18	0.00	**−1.56**
Gold Daily Price Change	−0.05	−0.15	**−0.27**



(a) Credit Card Fraud Detection

(b) MNIST

(c) Boston Housing

(d) Gold Daily Price Change

Figure 2.5: Regret bound analysis of the three stopping criteria across the four benchmark problems. For each criterion, the regret distribution of the 50 best-performing parameter settings is shown. The CDSC achieves lower mean regret and lower variance than both the maximum number of epochs and early stopping criteria in all cases.

Parameter sensitivity analysis revealed that the correlation threshold μ_{cdsc} is the dominant parameter of the CDSC, with random forest-based importance values of at least 0.70 across all datasets (Table 2.7). The window size ω and patience λ_{cdsc} were comparatively less important. Well-formulated problems with regular training cycles favoured positive μ_{cdsc} values ($\mu_{\text{cdsc}} = 0.75$ for Credit Card Fraud Detection, $\mu_{\text{cdsc}} = 0.40$ for MNIST); problems with fewer samples or irregular training dynamics benefited from negative thresholds ($\mu_{\text{cdsc}} < 0$ for Boston Housing and Gold).

Table 2.7: Parameter importance values of the CDSC, extracted from fitted random forest models, together with test accuracies of the random forest classifiers. The correlation threshold μ_{cdsc} dominates on every dataset (importance ≥ 0.70), whereas the window size and patience contribute comparatively little.

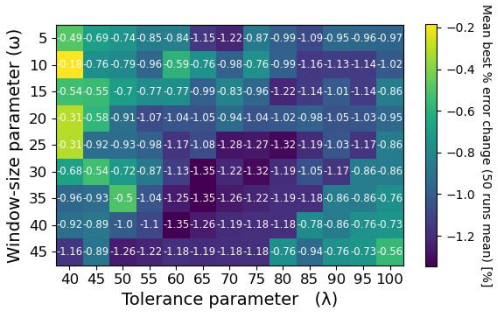
Dataset	Parameter importance			RF test accuracy [%]
	ω	λ_{cdsc}	μ_{cdsc}	
Credit Card Fraud Det.	0.081	0.160	0.758	82.6
MNIST	0.115	0.067	0.820	89.8
Boston Housing	0.090	0.073	0.837	78.9
Gold Daily Price Change	0.137	0.164	0.700	75.6

To further investigate the roles of ω and λ_{cdsc}, these parameters were fixed at each combination, and μ_{cdsc} was tuned to achieve the best possible $\overline{\%e}_{\text{ts}}$. The resulting heatmaps are shown in Figure 2.6. A distinct region of favourable ω–λ_{cdsc} combinations (visible as a blue "trench") exists for the Credit Card Fraud Detection, MNIST, and Gold Daily Price Change datasets. For the Boston Housing dataset, this region is less pronounced.

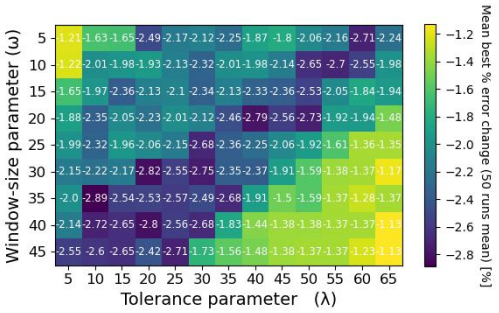
BO with GBRT converged to a suitable CDSC parameter setting in approximately 35–60 evaluations, depending on the problem and the size of the parameter space. Convergence statistics are reported in Table 2.8.

Table 2.8: BO (GBRT) convergence statistics: number of evaluations needed to find a sufficiently good CDSC parameter setting (50 runs per problem). BO located a competitive setting within a small fraction of the search space (a median of 18.5–37 evaluations against the 1800–4300 combinations that the parameter ranges yield when the threshold is discretised at 0.05 intervals); the median falls well below the mean because a few unfavourable initialisations lengthen the tail.

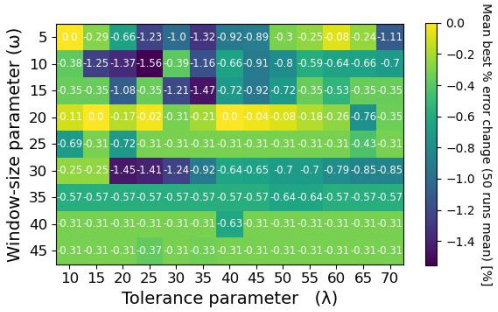
Dataset	Median	Mean	Std	No. param. combinations
Credit Card Fraud Det.	18.50	34.67	34.90	1800
MNIST	23.50	39.08	36.74	4300
Boston Housing	27.00	52.25	62.53	2700
Gold Daily Price Change	37.00	60.79	84.96	2700



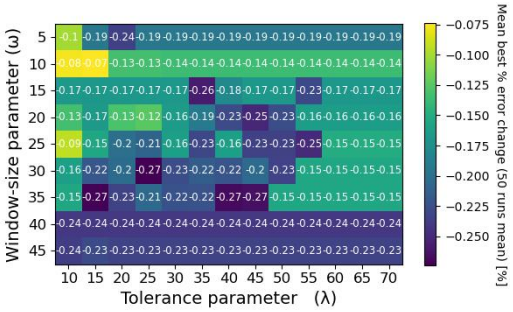
(a) Credit Card Fraud Detection



(b) MNIST



(c) Boston Housing



(d) Gold Daily Price Change

Figure 2.6: Best achievable $\overline{\%e}_{\text{ts}}$ with optimal μ_{cdsc} for each fixed λ_{cdsc}–ω combination. Darker blue regions indicate parameter regions where the CDSC achieves the greatest improvement. A connected trench of favourable ω–λ_{cdsc} combinations is visible for Credit Card Fraud Detection, MNIST, and Gold, but is far less pronounced for Boston Housing.

2.2.5 Discussion

The key advantage of the CDSC over existing stopping criteria is its ability to detect the onset of overfitting by monitoring the joint dynamics of the training and validation losses, rather than the validation loss alone. This makes it particularly effective in scenarios where the validation loss plateau is ambiguous or where training dynamics are noisy. When the validation loss exhibits an extended plateau (as is common with small datasets such as Boston Housing), early stopping terminates too soon regardless of the patience parameter, which is a fundamental limitation of validation-loss-only methods. Compared to the maximum number of epochs strategy, early stopping was only competitive on the Credit Card Fraud Detection dataset in the present benchmark; in practice, however, and without prior knowledge of the parameter-space behaviour, early stopping is easier to parameterise.

The regret-bound analysis shows that the CDSC not only achieves a better best-case result but also exhibits greater stability across near-optimal parameter configurations. Parameter sensitivity analysis confirms that the correlation threshold μ_{cdsc} is the dominant parameter; when the training process is noisy or the dataset is small, the rolling correlation naturally stays lower, requiring a negative threshold to allow training to continue for a sufficient number of epochs. The structured sensitivity patterns observed in the ω–λ_{cdsc} heatmaps correlate with well-formulated problems, well-designed models, and regular training processes, and they indicate that automatic parameter-tuning methods can be applied effectively. Although CDSC has a larger parameter space than early stopping, these structured patterns enable BO to find suitable parameters within a practical number of evaluations (35–60). The gap between the median and mean number of evaluations indicates that a small number of outlier runs, in which unfavourable initial evaluations lead to slower convergence, pull the mean upward; applying the sensitivity analysis to reduce the search space accelerates convergence further.

The method has limitations: its three-parameter search space is larger than that of early stopping, and its performance is sensitive to the correlation threshold μ_{cdsc}. However, the parameter sensitivity analysis provides practical guidance: the threshold μ_{cdsc} should be tuned first (starting from high positive values), and BO can efficiently automate the search. The method works best with cleaner datasets containing a large number of samples and models exhibiting regular training cycles. Larger batch sizes improve the stability of the rolling correlation curve, enhancing the criterion's precision. If the stop epoch produced by the CDSC equals N, this indicates that the criterion did not trigger within the allocated epoch budget and the learning process should be restarted with a larger N.

Based on the experimental results, the following practical guidelines for applying the CDSC are recommended:

- The CDSC works with automatic hyperparameter tuning methods; however, the size of the parameter space should be reduced to achieve faster convergence.
- Use datasets with more samples when possible, as they produce more stable correlation curves.
- Use larger batch sizes during training, as they improve the stability of the training process and, consequently, the rolling correlation signal.
- Start tuning μ_{cdsc} at high positive values (around 0.8–0.9) and decrease progressively while monitoring the loss function.
- The Pearson rolling correlation between the training and the validation losses serves as a diagnostic tool: noisy correlation curves and low initial correlation values indicate that the problem formulation, the model hyperparameters, or the training configuration should be revised.

The generality of CDSC, demonstrated across classification, regression, and time-series tasks using FFNN, CNN, and LSTM architectures, makes it applicable beyond the energy domain. Its application to EPF, which closes the methodological loop with the downstream contributions of this section, is presented in Section 2.3.

2.3 LSTM-based day-ahead electricity price forecasting with CDSC-enhanced training

2.3.1 Motivation and context

Accurate EPF for day-ahead markets is a prerequisite for the profitable operation of battery-enhanced renewable power plants, as described in Chapter 3. LSTM networks [25] have demonstrated strong performance on non-stationary time series, including electricity prices. However, the design of an LSTM network for this task involves a large hyperparameter space, lookback window size, network topology, activation functions, data normalisation, and regularisation strategy, and the literature lacked a systematic exploration of these design choices for day-ahead EPF on the Hungarian Power Exchange (HUPX).

Furthermore, the stopping criterion used during training directly affects the generalisation performance of the resulting model. The CDSC presented in Section 2.2 provides a principled mechanism to prevent overfitting during LSTM training, and its integration into the price-forecasting workflow is a key contribution. Extending the evaluation beyond HUPX to three structurally distinct DAMs with different temporal resolutions and volatility profiles (Section 2.3.3) assesses whether both the LSTM design decisions and the CDSC-enhanced training procedure generalise beyond a single market.

2.3.2 Method

The initial design targeting the HUPX DAM adopted a three-layer recurrent architecture: an input layer containing LSTM neurons, a hidden layer of regular (densely connected) neurons, and an output layer of 24 regular neurons, one per predicted hour. The model was trained end-to-end to minimise the mean absolute error (MAE) between predicted and realised day-ahead prices, with the stopping criterion held fixed at a maximum number of epochs.

The design was subsequently generalised to a stacked configuration with $D \in \{1, 2\}$ layers, hidden dimension $d_h \in [32, 256]$, and optional bidirectional processing. In this generalised configuration the full architecture is treated as part of the hyperparameter search space and optimised jointly with the training configuration and stopping-criterion parameters. The model is formulated as a direct multi-step forecaster: a single linear projection maps the final hidden state to n simultaneous output predictions, avoiding the recursive error accumulation inherent in iterative forecasting strategies [29, 30].

2.3.3 Experimental methodology

The HUPX study used hourly electricity prices from the HUPX DAM, covering the period 2010–2018. The first seven years (2010–2016) were used for training, 2017 for validation and hyperparameter selection, and 2018 for out-of-sample testing. The design process proceeded in two phases. In the first phase, the interaction between data normalisation and activation function was investigated, as these two hyperparameters are strongly correlated: for bounded activations such as sigmoid and tanh, the output range must match the normalised data range, whereas a linear activation imposes no such ceiling. A baseline network with 72 LSTM input neurons, 72 hidden neurons, and a 72-hour lookback was used. Three activation functions (linear, sigmoid, tanh) were tested with multiple normalisation ranges. The tested normalisation ranges and the output activation functions paired to them are illustrated in Figure 2.7.

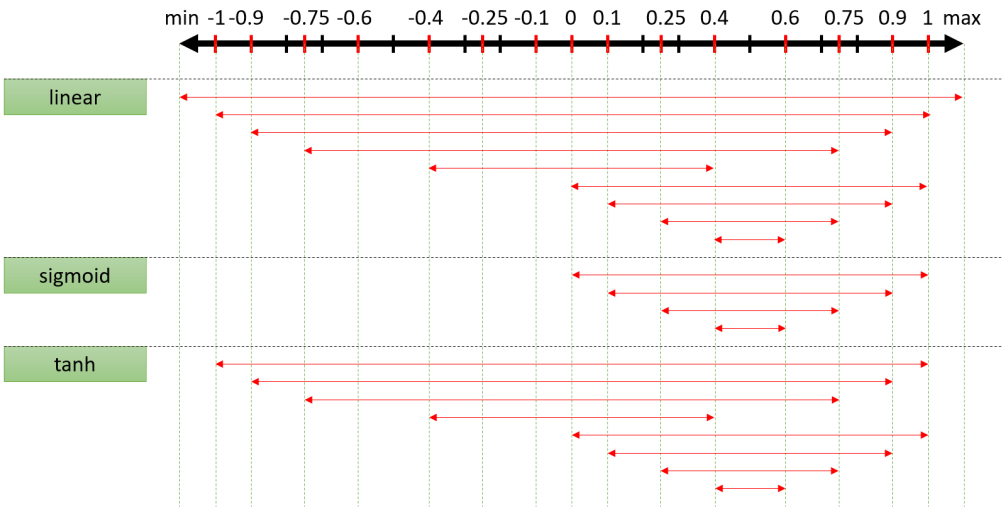


Figure 2.7: Normalisation ranges tested jointly with the three activation functions (linear, sigmoid, tanh). The bounds of the normalisation range were designed to match the output range of each activation function.

In the second phase, a full grid search was conducted over three hyperparameters: lookback window size $L \in \{6, 12, 24, 48, 72, 168\}$ hours, number of neurons per layer $\in \{6, 12, 24, 48, 72, 168\}$, and regularisation method (L1, L2, Elastic Net, Dropout, and unregularised). The hyperparameter settings are summarised in Table 2.9.

Table 2.9: Hyperparameters and their tested values for the LSTM grid search on HUPX. The second-phase grid pairs each lookback window and layer width with each regularisation method, holding the linear activation and $[0, 1]$ normalisation fixed from the first phase.

Hyperparameter	Tested values
Number of input data (lookback, hours)	6, 12, 24, 48, 72, 168
Number of neurons in input and hidden layers	6, 12, 24, 48, 72, 168
Regularisation method	none (unregularised), L1 (Lasso), L2 (Ridge), Elastic Net, Dropout

The evaluation was extended beyond HUPX to three structurally distinct DAMs, which also serve as the datasets for the foundation model study of Section 2.4:

1. **Germany (GER) DAM**: hourly settlement prices in EUR/MWh from the EPEX SPOT exchange, covering 2020–2024 (temporal resolution $\Delta t = 60$ min). GER, as a highly interconnected European exchange, exhibits comparatively moderate volatility.
2. **Singapore (SP) DAM**: half-hourly settlement prices in SGD/MWh from NEMS, covering 2020–2024 ($\Delta t = 30$ min). SP is less interconnected and exhibits more extreme price events, making forecasting more challenging.
3. **United Kingdom (UK) DAM**: hourly settlement prices in GBP/MWh from Nord Pool's N2EX, covering 2020–2024 ($\Delta t = 60$ min). UK prices are relatively smooth and liquid but still exhibit renewable-driven negative-price episodes.

Tables 2.10 to 2.12 present the descriptive statistics for the three markets. GER experienced substantial price volatility during 2021–2022 due to the European energy crisis, with mean prices rising from 30.4 EUR/MWh in 2020 to 236.5 EUR/MWh in 2022 before partially recovering to 77.8 EUR/MWh in 2024. SP showed a similar pattern, with mean prices peaking at 292.9 SGD/MWh in 2022, and UK peaked at 240.5 GBP/MWh in the same year. The year 2024 served as the primary held-out test set, excluded from all model selection procedures, while the year 2025 was retained for all three markets as a later held-out robustness sample used only for the out-of-year evaluations of Sections 2.3.4 and 2.4.4.

Table 2.10: Descriptive statistics of the GER price series (EUR/MWh). The mean price rose from 30.4 in 2020 to 236.5 at the 2022 energy-crisis peak before easing to 77.8 in 2024, and negative prices occur in every year.

Year	Mean	Median	Min	Max	Std
2020	30.41	30.82	−77.11	158.80	17.00
2021	97.55	75.48	−67.00	689.91	73.82
2022	236.48	208.19	−17.06	867.85	141.18
2023	95.11	97.48	−316.40	363.69	45.77
2024	77.80	78.44	−110.53	808.87	49.68
2025	89.48	92.95	−205.08	426.29	50.81

Table 2.11: Descriptive statistics of the SP price series (SGD/MWh). Prices are far more volatile than on the European markets, with the mean peaking at 292.9 in 2022 and half-hourly maxima between 4477 and 4545 in every year from 2021.

Year	Mean	Median	Min	Max	Std
2020	70.25	64.81	0.23	1241.07	47.32
2021	195.62	104.36	−4.84	4534.10	363.31
2022	292.91	236.25	56.91	4542.07	269.05
2023	246.72	158.93	74.25	4544.67	340.82
2024	163.12	121.05	−18.95	4527.25	247.43
2025	117.87	101.09	50.21	4477.20	111.18

Table 2.12: Descriptive statistics of the UK price series (GBP/MWh). The mean peaked at 240.5 in 2022 and fell to 85.8 by 2024, with renewable-driven negative prices in every year.

Year	Mean	Median	Min	Max	Std
2020	39.59	38.03	−43.32	387.67	18.49
2021	137.64	103.14	−16.92	2575.91	119.79
2022	240.53	214.49	−50.55	2165.04	130.57
2023	108.00	104.54	−60.19	292.56	40.82
2024	85.81	86.87	−27.83	583.16	35.57
2025	94.36	94.30	−38.38	1061.56	42.78

Two prediction horizons were evaluated, $n = 4$ and $n = 40$, yielding six prediction scenarios. A joint hyperparameter search space was defined covering model architecture (number of layers, hidden size, bidirectional mode) and training configuration (batch size, learning rate, gradient clipping norm, loss function, learning rate scheduler, λ_{L1} and λ_{L2} regularisation weights). The lookback window length $L \in [16, 2048]$ timesteps and the stopping-criterion parameters themselves were included in the same search space. For the CDSC, the window size $\omega \in [1, 20]$, correlation threshold $\mu_{\text{cdsc}} \in [0.0, 0.9]$, and patience $\lambda_{\text{cdsc}} \in [1, 30]$ were optimised; for early stopping, the patience $p_{\text{es}} \in [5, 30]$ was optimised analogously. This design ensures a fair comparison: both stopping criteria are given equal opportunity to find their best operating point within the joint optimisation framework.

Hyperparameter optimisation was performed using the TPE [96] within the Optuna framework [101] on an HPC cluster. Each trial was repeated with $R = 3$ independent random seeds, and the optimisation objective was defined as the mean validation MAE averaged across all CV folds and repetitions.

Two temporal validation strategies were evaluated, illustrated in Figure 2.8: (i) a single-fold holdout design in which 2020–2021 and the first part of 2022 form the training set while the remainder of 2022 together with 2023 form the validation set, and (ii) a four-fold leave-one-year-out CV scheme over the years 2020–2023. Both strategies apply an embargo zone of $L + n$ observations at each partition boundary to prevent information leakage [45]; the embargo width also fixes the exact split point within 2022 for the single-fold design. In both cases, the year 2024 served as the primary held-out test set, excluded from all model selection procedures, while the year 2025 was retained as a later held-out robustness sample, likewise excluded from all tuning and model selection and used only for the out-of-year evaluation reported below.

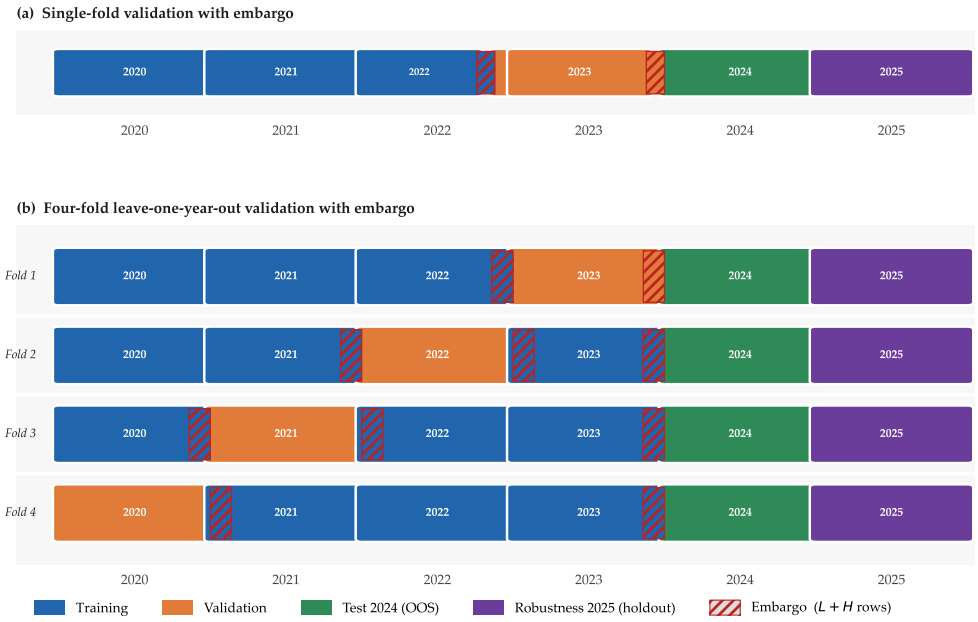


Figure 2.8: The two temporal validation strategies and the two held-out evaluation years. Panel (a): single-fold validation, with 2020–2021 and the first part of 2022 as the training set and the remainder of 2022 together with 2023 as the validation set. Panel (b): four-fold leave-one-year-out CV over 2020–2023, each year serving once as the validation fold while the remaining three form the training set. Hatched regions mark the embargo zones of width $L + n$ inserted at every partition boundary to prevent leakage from lagged features. Both the 2024 test year and the 2025 robustness year sit outside every tuning and model-selection step, which is what makes the out-of-year comparison of Section 2.3.4 a genuine holdout.

2.3.4 Results

The results of the normalisation–activation function investigation on HUPX are presented in Table 2.13. The best-performing combination was the linear activation function with input data normalised to the range $[0, 1]$, achieving MAE = 8.03 EUR/MWh. Without normalisation, the same architecture achieved MAE = 9.00 EUR/MWh. The best tanh-based network scored MAE = 8.14 EUR/MWh and was normalised to $[0, 1]$. Sigmoid-based tests performed best with normalisation to $[0.1, 0.9]$ (MAE = 8.99 EUR/MWh).

Table 2.13: Mean absolute errors (EUR/MWh) for different activation functions and normalisation range combinations. The lowest error (MAE = 8.03) was obtained with linear activation and $[0, 1]$ normalisation.

Activation function	Normalisation range			
	$[0, 1]$	$[0.1, 0.9]$	$[0.25, 0.75]$	$[0.4, 0.6]$
sigmoid	9.11	8.99	9.16	10.27
tanh	8.14	8.40	8.82	10.71
linear	**8.03**	8.30	8.55	10.49
Activation function	Normalisation range			
	$[-1, 1]$	$[-0.9, 0.9]$	$[-0.75, 0.75]$	$[-0.4, 0.4]$
sigmoid	–	–	–	–
tanh	8.44	8.69	8.58	8.63
linear	8.19	9.06	8.67	8.66

The full grid search over lookback window size, number of neurons per layer, and regularisation method identified the best configuration: $L = 48$ hours, 168 LSTM neurons in the input layer, 168 regular neurons in the hidden layer, and no regularisation.

The effect of regularisation on the best-performing network topology is summarised in Table 2.14. The unregularised network consistently outperformed all regularised variants. Mild Dropout ($p_{\text{drop}} = 0.01$) increased MAE by 0.57 EUR/MWh, while L1 regularisation ($\lambda_{L1} = 0.01$) increased it by 0.68 EUR/MWh.

Table 2.14: Effect of regularisation methods on the best-performing LSTM network topology (168 LSTM neurons, 48-hour lookback, linear activation, $[0, 1]$ normalisation). The unregularised network is the most accurate (MAE = 7.48 EUR/MWh); every explicit weight or dropout penalty raises the error, indicating that the LSTM gates already provide sufficient regularisation on this task.

Regularisation method	Parameters	MAE (EUR/MWh)
None (unregularised)	–	**7.48**
L1 (Lasso)	$\lambda_{L1} = 0.01$	8.16
L2 (Ridge)	$\lambda_{L2} = 0.01$	8.10
Elastic Net	$\lambda_{L1} = \lambda_{L2} = 0.01$	8.27
Dropout	$p_{\text{drop}} = 0.01$	8.05
Dropout	$p_{\text{drop}} = 0.05$	8.07

The final HUPX LSTM network consisted of 168 LSTM neurons (input layer), 168 regular neurons (hidden layer), and 24 regular neurons (output layer). The input comprised 48 previous hourly electricity prices, normalised to $[0, 1]$, with linear activation in both hidden and output layers. The model was trained for 100 epochs with a batch size of 168. The best configuration identified by the grid search achieved MAE = 7.48 EUR/MWh.

The extended LSTM results across GER, SP, and UK are reported in Table 2.15 across all combinations of CV strategy and stopping method. All MAE values are computed on the held-out 2024 test set. The CDSC attained a lower test MAE than standard early stopping in 10 of the 12 matched configurations, with the advantage concentrated at the longer prediction horizon.

Table 2.15: LSTM results across datasets, prediction horizons, CV types, and stopping methods. All MAE values are computed on the held-out 2024 test set. Results are ordered by test MAE within each scenario (best first). The CDSC attained a lower test MAE than early stopping in 10 of the 12 matched configurations, with the advantage concentrated at the longer prediction horizon.

Dataset	n	CV type	Stopping method	L	Test MAE
GER	4	CV	CDSC	144	15.35
GER	4	CV	early stopping	28	15.70
GER	4	no CV	CDSC	34	15.95
GER	4	no CV	early stopping	28	16.28
GER	40	no CV	CDSC	167	28.80
GER	40	CV	CDSC	163	30.21
GER	40	no CV	early stopping	174	31.55
GER	40	CV	early stopping	241	31.92
SP	4	CV	CDSC	205	42.92
SP	4	no CV	early stopping	52	44.33
SP	4	CV	early stopping	86	46.57
SP	4	no CV	CDSC	300	46.69
SP	40	no CV	CDSC	313	53.15
SP	40	CV	CDSC	329	54.66
SP	40	CV	early stopping	71	57.34
SP	40	no CV	early stopping	158	63.81
UK	4	no CV	early stopping	182	12.22
UK	4	no CV	CDSC	36	12.61
UK	4	CV	CDSC	166	13.39
UK	4	CV	early stopping	223	15.21
UK	40	CV	CDSC	74	19.21
UK	40	CV	early stopping	26	19.65
UK	40	no CV	CDSC	57	19.71
UK	40	no CV	early stopping	22	21.38

On GER the CDSC achieved a lower test MAE than early stopping across all four configurations, with the largest advantage at the longer prediction horizon ($n = 40$) without CV: MAE = 28.80 for the CDSC versus 31.55 for early stopping, a reduction of 8.7%. For $n = 4$ with CV, the CDSC achieved a test MAE of 15.35 compared to 15.70 for early stopping. On SP the pattern was scenario-dependent: for $n = 4$ without CV early stopping slightly outperformed the CDSC (44.33 versus 46.69), whereas for the same horizon with CV the CDSC yielded the best result (42.92 versus 46.57). For $n = 40$ the CDSC dominated in both validation strategies, with the widest margin without CV: MAE = 53.15 for the CDSC versus 63.81 for early stopping, a reduction of 16.7% that is the largest CDSC advantage across all scenarios, where the early-stopping value of 63.81 was the worst LSTM configuration overall. On UK the two criteria were close: the CDSC was the more accurate in three of the four configurations (13.39 versus 15.21 and 19.21 versus 19.65 under CV, and 19.71 versus 21.38 for $n = 40$ without CV), while early stopping was marginally better for $n = 4$ without CV (12.22 versus 12.61). Overall the CDSC attained the lower test MAE in 10 of the 12 matched configurations; the two exceptions were the SP and UK short-horizon ($n = 4$) single-fold settings, where standard early stopping remained competitive.

A weak tendency was observed for the four-fold CV design to select longer lookback windows than single-fold validation, consistent with the multi-fold scheme exposing the optimiser to several annual periodicities; the selected lookback windows nonetheless overlapped substantially between the two strategies (Table 2.15).

The optimised stopping-criterion parameters are reported in Table 2.16. CDSC window sizes ranged from $\omega = 10$ to $\omega = 20$, correlation thresholds from $\mu_{\text{cdsc}} = 0.0$ to $\mu_{\text{cdsc}} = 0.7$, and patience values from $\lambda_{\text{cdsc}} = 5$ to $\lambda_{\text{cdsc}} = 30$.

Table 2.16: Optimised stopping criterion parameters for LSTM models across all configurations. The CDSC settings vary widely across markets and horizons (threshold μ_{cdsc} from 0.0 to 0.7, window from 10 to 20, patience from 5 to 30), which is why they were tuned jointly with the architecture rather than fixed in advance.

Dataset	n	CV type	Stopping	p_{es}	ω	μ_{cdsc}	λ_{cdsc}
GER	4	CV	CDSC	–	10	0.4	25
GER	4	CV	early stopping	20	–	–	–
GER	4	no CV	CDSC	–	20	0.3	15
GER	4	no CV	early stopping	25	–	–	–
GER	40	no CV	CDSC	–	10	0.7	20
GER	40	CV	CDSC	–	10	0.5	30
GER	40	no CV	early stopping	10	–	–	–
GER	40	CV	early stopping	25	–	–	–
SP	4	CV	CDSC	–	15	0.4	5
SP	4	no CV	early stopping	25	–	–	–
SP	4	CV	early stopping	25	–	–	–
SP	4	no CV	CDSC	–	20	0.2	5
SP	40	no CV	CDSC	–	20	0.0	15
SP	40	CV	CDSC	–	15	0.1	30
SP	40	CV	early stopping	20	–	–	–
SP	40	no CV	early stopping	5	–	–	–
UK	4	no CV	early stopping	15	–	–	–
UK	4	no CV	CDSC	–	10	0.1	25
UK	4	CV	CDSC	–	20	0.4	20
UK	4	CV	early stopping	15	–	–	–
UK	40	CV	CDSC	–	10	0.0	20
UK	40	CV	early stopping	15	–	–	–
UK	40	no CV	CDSC	–	15	0.2	20
UK	40	no CV	early stopping	5	–	–	–

The statistical significance of the stopping-criterion differences was assessed with a two-sided Diebold–Mariano (DM) test [24] applied to the per-origin held-out losses, using the small-sample correction of Harvey, Leybourne, and Newbold [102]. Table 2.17 reports the loss differential Δ MAE, defined as the CDSC test MAE minus the early-stopping test MAE, so that a negative value marks a lower error for the CDSC. Six of the 12 differences were statistically significant on the 2024 test set: the CDSC was significantly more accurate at $n = 40$ on both GER and SP under both validation modes and on SP at $n = 4$ with CV, the largest margin being Δ MAE $= -10.66$ on SP at $n = 40$ without CV, whereas early stopping was significantly more accurate only on SP at $n = 4$ without CV. At $n = 4$ on GER and at both horizons on UK the two criteria were statistically indistinguishable. The ordering was broadly preserved on the 2025 holdout, with more differences reaching significance on GER and SP, and the UK criteria remaining indistinguishable.

Table 2.17: LSTM stopping-criterion comparison (CDSC versus early stopping) by a two-sided DM test with the Harvey–Leybourne–Newbold small-sample correction, at the 5% significance level. Δ MAE is the CDSC test MAE minus the early-stopping test MAE; negative values indicate a lower error for the CDSC. "Sig." marks whether the difference is statistically significant. Results are shown for the 2024 test year and the 2025 holdout year.

			2024		2025	
Dataset	n	CV type	Δ MAE	Sig.	Δ MAE	Sig.
GER	4	no CV	−0.33	No	+0.47	Yes
GER	4	CV	−0.35	No	+0.77	Yes
GER	40	no CV	−2.75	Yes	−1.64	Yes
GER	40	CV	−1.71	Yes	−1.96	Yes
SP	4	no CV	+2.36	Yes	+3.65	Yes
SP	4	CV	−3.65	Yes	−2.52	Yes
SP	40	no CV	−10.66	Yes	−9.73	Yes
SP	40	CV	−2.68	Yes	−1.12	Yes
UK	4	no CV	+0.39	No	+0.21	No
UK	4	CV	−1.82	No	+0.17	No
UK	40	no CV	−1.67	No	−0.01	No
UK	40	CV	−0.44	No	+0.05	No

Robustness on the 2025 holdout

The relative ordering of the two stopping methods on the 2024 test year largely carried over to the held-out 2025 year, as Table 2.18 shows. The CDSC retained the lower 2025 MAE at $n = 40$ on GER and SP and at $n = 4$ on SP, while early stopping retained its 2024 lead at $n = 4$ on UK. The within-variant ranking reversed in two scenarios: at $n = 4$ on GER, where early stopping overtook the CDSC in a statistically significant change, and at $n = 40$ on UK, where the reversal was statistically indistinguishable. Market-level holdout drift mirrored the pattern seen for the other model families: SP showed large negative drift, indicating that 2025 was easier to forecast (its $n = 4$ CDSC error fell from 42.92 to 20.02 and its $n = 40$ error from 53.15 to 26.79); GER showed positive drift (its $n = 4$ CDSC error rose from 15.35 to 19.09); and UK showed only small positive drift. Despite this stable ordering, neither LSTM variant entered the 2025 Model Confidence Set (MCS) [103] in any of the six scenarios, as Table 2.19 records: at the 10% significance level (a 90% confidence level) both configurations were excluded as significantly less accurate than the leading model, so the optimised LSTM remained below the leading tier on the unseen year in every market and horizon.

Table 2.18: LSTM test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; the year 2025 was not seen during tuning. **Bold** marks the lower MAE per column within the LSTM family. The relative ranking of the two stopping methods largely carries over from 2024 to the unseen 2025 year.

	GER-4		GER-40		SP-4		SP-40		UK-4		UK-40	
Model	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025
CDSC	**15.35**	19.09	**28.80**	**33.10**	**42.92**	**20.02**	**53.15**	**26.79**	12.61	14.76	**19.21**	23.66
Early stopping	15.70	**18.32**	31.55	34.74	44.33	22.93	57.34	27.63	**12.22**	**14.55**	19.65	**23.62**

Table 2.19: Membership of the LSTM models in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Neither stopping criterion entered the set in any scenario.

Model	GER-4	GER-40	SP-4	SP-40	UK-4	UK-40
CDSC	No	No	No	No	No	No
Early stopping	No	No	No	No	No	No

2.3.5 Discussion

The systematic HUPX grid search revealed several practical insights. The 48-hour (two-day) lookback window provides sufficient temporal context without introducing noise from older, less relevant patterns on HUPX. The linear activation function outperformed sigmoid and tanh, likely because electricity prices are not bounded and a linear function imposes no output ceiling. The linear-activated network with $[0, 1]$ normalisation trained well and showed no signs of overfitting, motivating its selection for further investigation. The finding that weight-based regularisation degrades LSTM performance on this task is particularly noteworthy: regularisation is generally considered beneficial for NNs, but on this task no explicit weight or dropout penalty improved on the unregularised network. A plausible explanation is that the gating mechanism already supplies sufficient implicit regularisation, though isolating that mechanism would require a dedicated ablation. This finding further

motivates the use of the CDSC as an alternative regularisation mechanism that does not interfere with the weight structure.

In the downstream trading application presented in Chapter 3, the predicted HUPX prices lowered the profit increase over the immediate-dispatch baseline by only 1.93–6.93 percentage points relative to the oracle (perfect-information) scenario, validating the practical utility of the LSTM predictions for battery-enhanced trading strategies.

The extended evaluation on GER, SP, and UK confirmed and generalised the insights from the HUPX study. The CDSC attained a lower test MAE than early stopping in 10 of the 12 matched LSTM configurations, and the DM tests showed that its advantage was horizon-dependent: the gains were concentrated, and statistically significant, at the longer prediction horizon ($n = 40$) and on the high-volatility SP, whereas at the short horizon ($n = 4$) on the smoother GER and UK the two criteria were statistically indistinguishable and early stopping was occasionally preferable. The mechanism underlying the CDSC advantage is that it monitors the joint trajectory of training and validation losses, making it more robust to the plateaus and temporary reversals that characterise loss curves in noisy, spike-prone electricity price series; standard patience-based early stopping proved competitive mainly where the validation signal was relatively clean, that is, at short horizons on the less volatile markets.

The four-fold leave-one-year-out CV scheme improved LSTM generalisation in many configurations, though the effect was neither universal nor uniform. The clearest and most consistent benefits appeared on the volatile SP, where annual price distributions varied dramatically during the 2020–2023 period; by exposing the hyperparameter optimiser to diverse annual market conditions, the multi-fold design selected configurations that transferred more reliably to the unseen 2024 test year. On the smoother GER at $n = 40$ and on UK at $n = 4$, however, single-fold validation was preferable, confirming that the value of CV is market- and horizon-specific. The weak tendency for CV to favour longer lookback windows than single-fold validation is consistent with the multi-fold design encouraging the optimiser to select models that exploit longer-range temporal dependencies.

The scenario-specific nature of the optimised CDSC parameters is informative. A threshold of $\mu_{\text{cdsc}} = 0.0$ (SP, $n = 40$, no CV) indicates that any decorrelation between the training and validation trajectories was treated as a stopping signal; conversely, a threshold of $\mu_{\text{cdsc}} = 0.7$ (GER, $n = 40$, no CV) tolerated substantial decorrelation before triggering termination. This range echoes the finding that μ_{cdsc} is the dominant parameter of the CDSC and confirms that its optimal value is market- and horizon-specific, requiring the kind of automated optimisation framework employed here. This wide range confirms the finding that the CDSC parameters interact with the characteristics of the dataset and task in ways that cannot be anticipated a priori, and validates the decision to include them in the joint hyperparameter search rather than fixing them based on the general benchmark results of Section 2.2.4.

Taken together, results demonstrate that the CDSC is a viable and often superior alternative to standard early stopping for LSTM-based day-ahead EPF, and that the joint optimisation of stopping-criterion parameters alongside model architecture yields configurations that would not be discovered by treating these as fixed design choices. The following section extends the evaluation to Transformer-based foundation models and benchmarks them against the best LSTM configurations established here.

2.4 Transformer-based foundation models for direct multi-step electricity price forecasting

2.4.1 Motivation and context

The LSTM-based approach developed in Section 2.3 demonstrated that RNNs, when combined with systematic hyperparameter optimisation and the CDSC, can achieve competitive forecasting accuracy on day-ahead electricity prices. However, recurrent architectures process input sequences step by step, which limits their ability to capture long-range dependencies efficiently and precludes parallelisation during training. Transformer-based architectures [104], which employ self-attention mechanisms to model dependencies across all positions in the input sequence simultaneously, have been proposed as an alternative paradigm for time-series forecasting. Several Transformer variants have been developed specifically for temporal data. The Informer was proposed as an efficient Transformer for long-sequence time-series forecasting [32]. The Autoformer combines series decomposition with an auto-correlation mechanism for long-term forecasting [33]. The Temporal Fusion Transformer targets interpretable multi-horizon time-series forecasting [34]. PatchTST is a patch-based Transformer for long-term time-series forecasting [35]. Zeng et al. examined whether Transformers are effective for time-series forecasting, reporting that their advantage over simpler models can diminish on moderate-sized datasets [36].

Most recently, a new paradigm has emerged: pre-trained time-series foundation models trained on large, heterogeneous corpora of time series from diverse domains, then applied to new tasks in a zero-shot regime without any task-specific gradient updates. These models draw on the implicit cross-domain knowledge encoded during large-scale pre-training, requiring only the selection of an appropriate context window length at inference time. Early benchmarks on EPF tasks indicated that foundation models perform on par with traditional approaches but did not consistently outperform well-tuned task-specific models [41]. However, the sensitivity of these foundation models to the context window length had not been systematically optimised on day-ahead electricity market datasets, and several models evaluated by Hornek et al. had been superseded by newer versions by the time this evaluation was carried out.

This section addresses this gap by evaluating six state-of-the-art Transformer-based foundation models with systematic context-length optimisation, and comparing their zero-shot forecasting performance against the optimised LSTM models from Section 2.3. The evaluation spans three structurally distinct DAMs and two prediction horizons, enabling a controlled assessment of cross-market generalisability and the relative merits of supervised recurrent models versus zero-shot foundation models for EPF.

2.4.2 Method

Six pre-trained time-series foundation models from three model families were evaluated in a zero-shot inference regime. No gradient updates were performed on the target datasets; the models were used exactly as released by their developers. Table 2.20 summarises the models and their characteristics.

Table 2.20: Transformer-based foundation models evaluated in zero-shot inference. Model sizes are reported as in the original publications. The six models are drawn from three families (Chronos, TimesFM, Moirai) and span 9M to 205M parameters, all applied without any task-specific gradient updates.

Model	Developer	Architecture	Size	Reference
Chronos-Bolt (Tiny)	Amazon	Enc–dec Transformer	9M	[38]
Chronos-Bolt (Small)	Amazon	Enc–dec Transformer	48M	[38]
Chronos-Bolt (Base)	Amazon	Enc–dec Transformer	205M	[38]
Chronos 2	Amazon	Enc–dec Transformer	120M	[37, 105]
TimesFM 2.5	Google	Decoder-only	200M	[39, 106]
Moirai-2 (Small)	Salesforce	Masked encoder	11.4M	[40, 107]

Chronos (Amazon). Chronos [37] treats time-series forecasting as a sequence-to-sequence token generation problem. Continuous values are normalised via per-series scaling, tokenised into a fixed vocabulary of integer bins, and modelled with a cross-entropy training objective. Point forecasts are obtained as the median of the predictive distribution. Three variants from the Chronos-Bolt family [38], which use an efficient encoder–decoder architecture with improved training, were evaluated alongside Chronos 2, the latest generation model.

TimesFM (Google). TimesFM 2.5 [39] is a 200-million-parameter decoder-only Transformer pre-trained on a large corpus of real-world and synthetic time series. Unlike Chronos, TimesFM operates directly on continuous values rather than tokenised inputs, using patching to aggregate multiple time steps into a single input token. The model produces point forecasts directly.

Moirai (Salesforce). Moirai [40] is a universal time-series forecasting Transformer trained on LOTSA, a large heterogeneous corpus of publicly available time-series datasets. It employs a masked encoder architecture with frequency-aware patch embeddings and a mixture-of-distributions output head. Point estimates were obtained as the median of 100 Monte Carlo samples drawn from the predictive distribution. The Moirai-2 (Small) variant was evaluated by interfacing with the GluonTS [108] prediction API.

2.4.3 Experimental methodology

The foundation models were evaluated on the same three DAM price series used for the LSTM study and described in Section 2.3.3: GER, SP, and UK.

Two prediction horizons were evaluated: $n = 4$ and $n = 40$ timesteps. For the hourly GER series, these correspond to 4-hour and 40-hour-ahead forecasts; for the half-hourly SP series, to 2-hour and 20-hour-ahead forecasts. The combination of three markets and two horizons yields the six prediction scenarios that define the experimental grid.

In the zero-shot setting used here the foundation models require no task-specific training, but they remain sensitive to the length of the context window L provided at inference time. Rather than fixing L a priori, it was treated as a study-level hyperparameter and optimised via exhaustive grid search (Optuna GridSampler) over the discrete range $L \in \{24, 28, 32, \ldots, 2048\}$ (step 4, yielding 507 candidate values). A separate Optuna study was created per foundation model, enabling independent and comparable evaluation of each model's context-length sensitivity.

The optimisation objective was the mean validation MAE computed directly from zero-shot predictions over the validation partition. Because foundation models are deterministic given L (no stochastic training), repetitions were not required ($R = 1$), except for Moirai, where a fixed random seed was used when drawing the 100 samples from the predictive distribution. The selected optimal context length L^* was then used to produce the final test-set forecast on the 2024 data.

2.4.4 Results

Table 2.21 reports the test MAE of the six foundation models across all six prediction scenarios. Chronos 2 achieved the lowest test MAE in five of the six scenarios, the exception being UK at $n = 40$, where Moirai-2 (Small) was marginally lower.

A clear performance hierarchy emerged across models: Chronos 2 led in five of the six scenarios (Moirai-2 (Small) was marginally ahead on UK at $n = 40$), followed by a competitive second tier comprising Chronos-Bolt (Base and Small) and Moirai-2 (Small), typically within 1–2 MAE units of Chronos 2. TimesFM 2.5 and Chronos-Bolt (Tiny) consistently ranked lower. On GER with $n = 40$ the top four models were separated by only 0.52 MAE units.

The optimised context lengths exhibited a clear market-dependent pattern. GER models selected contexts of 204–520 timesteps (approximately 8.5–21.7 days at hourly resolution), SP models considerably longer contexts of 388–2036 timesteps (approximately 8.1–42.4 days at half-hourly resolution), and UK models the longest windows of 552–2048 timesteps (approximately 23–85 days at hourly resolution), with Chronos 2 reaching the 2048-step search-space limit at both horizons. When converted to calendar time the ranges are of a broadly similar order across markets, spanning roughly one week to a few months, and the UK optimum at the search boundary suggests that its true optimum may lie beyond the tested range.

Across all six prediction scenarios, the best foundation model outperformed the best LSTM configuration from Section 2.3.4. The comparison is summarised in Table 2.22. The test MAE improvements over the best LSTM were 40.1% and 30.0% on GER, 18.4% and 19.0% on SP, and 40.8% and 26.4% on UK, at $n = 4$ and $n = 40$ respectively. The best foundation model was Chronos 2 in every scenario except UK at $n = 40$, where Moirai-2 (Small) was marginally the most accurate.

Within the LSTM family, the spread between the best and worst configurations was 5.7% on GER at $n = 4$ and 16.7% on SP at $n = 40$.

Table 2.21: Foundation model results across datasets and prediction horizons. Models are ordered by test MAE within each scenario. Chronos 2 attained the lowest test MAE in five of the six scenarios, the exception being UK at $n = 40$, where Moirai-2 (Small) was marginally lower.

Dataset	n	Model	L^*	Val. MAE	Test MAE
GER	4	Chronos 2	356	8.76	9.19
GER	4	Moirai-2 (Small)	512	9.26	9.57
GER	4	Chronos-Bolt (Base)	368	9.10	9.77
GER	4	Chronos-Bolt (Small)	520	9.39	9.80
GER	4	TimesFM 2.5	208	9.73	10.45
GER	4	Chronos-Bolt (Tiny)	352	10.06	11.00
GER	40	Chronos 2	388	21.09	20.16
GER	40	Chronos-Bolt (Small)	508	21.79	20.24
GER	40	Chronos-Bolt (Base)	352	21.58	20.59
GER	40	Moirai-2 (Small)	380	21.99	20.68
GER	40	Chronos-Bolt (Tiny)	512	22.52	20.88
GER	40	TimesFM 2.5	204	22.91	22.27
SP	4	Chronos 2	844	58.68	35.03
SP	4	Chronos-Bolt (Base)	892	60.70	35.56
SP	4	TimesFM 2.5	808	61.46	36.01
SP	4	Moirai-2 (Small)	836	61.16	36.21
SP	4	Chronos-Bolt (Small)	864	60.66	36.28
SP	4	Chronos-Bolt (Tiny)	388	63.24	37.12
SP	40	Chronos 2	1544	81.74	43.03
SP	40	Moirai-2 (Small)	2036	83.60	44.15
SP	40	Chronos-Bolt (Base)	780	84.66	44.19
SP	40	TimesFM 2.5	1344	85.37	44.65
SP	40	Chronos-Bolt (Small)	764	86.09	44.74
SP	40	Chronos-Bolt (Tiny)	880	87.40	45.73
UK	4	Chronos 2	2048	7.70	7.23
UK	4	Moirai-2 (Small)	1876	8.14	7.66
UK	4	Chronos-Bolt (Base)	584	8.24	7.93
UK	4	Chronos-Bolt (Small)	552	8.44	8.10
UK	4	TimesFM 2.5	1136	8.59	8.17
UK	4	Chronos-Bolt (Tiny)	572	9.02	8.50
UK	40	Moirai-2 (Small)	1492	14.62	14.13
UK	40	Chronos 2	2048	13.95	14.19
UK	40	Chronos-Bolt (Tiny)	944	15.18	14.83
UK	40	Chronos-Bolt (Base)	624	14.50	14.96
UK	40	Chronos-Bolt (Small)	568	14.78	15.04
UK	40	TimesFM 2.5	1944	15.02	15.07

Table 2.22: Comparison of the best foundation model (FM) with the best LSTM configuration in each scenario. The best FM is Chronos 2 in every scenario except UK at $n = 40$, where Moirai-2 (Small) is marginally lower. CV: cross-validation; ES: early stopping. The improvement is computed as $(\mathrm{MAE}_{\mathrm{LSTM}} - \mathrm{MAE}_{\mathrm{FM}})/\mathrm{MAE}_{\mathrm{LSTM}} \times 100\%$.

Dataset	n	Best LSTM test MAE	Best FM test MAE	Improvement	Best LSTM config.
GER	4	15.35	9.19	40.1%	CV + CDSC
GER	40	28.80	20.16	30.0%	no CV + CDSC
SP	4	42.92	35.03	18.4%	CV + CDSC
SP	40	53.15	43.03	19.0%	no CV + CDSC
UK	4	12.22	7.23	40.8%	no CV + ES
UK	40	19.21	14.13	26.4%	CV + CDSC

Conventional baseline comparison

To place the neural models in context, three conventional EPF baselines were evaluated on the same six scenarios: a seasonal-naive persistence baseline (predicting each step from the price one week earlier), the Lasso Estimated AutoRegressive (LEAR) linear benchmark [23], and the Neural Basis Expansion Analysis for Time Series (N-BEATS) network [109]. The best baseline configurations are reported in Table 2.23. N-BEATS attained the lowest test MAE among the baselines in five of the six scenarios, LEAR being marginally more accurate only on GER at $n = 40$ (21.56 versus 22.20). The seasonal-naive baseline was the least accurate in every scenario and defined the accuracy floor, with test MAE near 30.9 on GER, 74.8 on SP, and 23.2 on UK, and its error almost unchanged between the two horizons, whereas LEAR and N-BEATS degraded at the longer horizon. Crucially, even the strongest conventional baseline remained above the best foundation model in every scenario, so the zero-shot foundation models set a higher accuracy standard than both the optimised LSTM and the established conventional benchmarks. The panel spans the standard statistical (LEAR), neural (N-BEATS), and persistence (seasonal-naive) benchmarks used in the EPF literature; tree-based gradient-boosting models (for example XGBoost or LightGBM) were not included and remain a natural further baseline.

The within-family ordering of the conventional baselines was largely stable across the two evaluation years, as Table 2.24 shows. N-BEATS attained the lower test MAE on SP and UK at both horizons and in both years, while LEAR was more accurate on GER at $n = 40$; at $n = 4$ on GER the two were close, with N-BEATS lower in 2024 and LEAR lower in 2025. The seasonal-naive baseline trailed throughout. Market-level holdout drift followed the pattern seen for the other model families: SP showed large negative drift, indicating that 2025 was the easier year, with the N-BEATS error at $n = 4$ falling from 39.11 to 18.46 and the LEAR error from 46.46 to 21.68; GER showed positive drift, the LEAR error at $n = 4$ rising from 11.15 to 12.90; and UK showed only small positive drift.

The conventional family remained largely outside the foundation-model-led 2025 MCS, as Table 2.25 shows. LEAR and the seasonal-naive baseline fell outside the 2025 set in every scenario, and N-BEATS entered it in only two cases, SP at $n = 4$ and UK at $n = 40$. The carry-over to the unseen year therefore preserved the relative ordering within the conventional baselines without bringing them level with the most accurate models on 2025.

Regime-segmented error analysis

To examine performance beyond aggregate accuracy, the 2024 test-year errors were segmented by price regime. A forecast point was labelled a spike when the ground-truth price exceeded the annual mean by more than two standard deviations, negative when the price was below zero, and normal otherwise; directional accuracy was measured as the fraction

Table 2.23: Best conventional baseline configurations (test MAE on the 2024 test set) across the six scenarios. For LEAR and N-BEATS the better of the single-fold and CV configurations is shown; the seasonal-naive baseline has no tunable parameters. L: lookback window (timesteps). N-BEATS was the most accurate baseline in five of the six scenarios, yet even the strongest baseline remained less accurate than the best foundation model in every scenario.

Dataset	n	Model	L	Test MAE
GER	4	N-BEATS	168	11.09
GER	4	LEAR	560	11.15
GER	4	Seasonal-naive	168	30.86
GER	40	LEAR	741	21.56
GER	40	N-BEATS	168	22.20
GER	40	Seasonal-naive	168	30.88
SP	4	N-BEATS	48	39.11
SP	4	LEAR	468	46.46
SP	4	Seasonal-naive	336	74.76
SP	40	N-BEATS	48	52.15
SP	40	LEAR	447	62.28
SP	40	Seasonal-naive	336	74.77
UK	4	N-BEATS	192	9.06
UK	4	LEAR	178	10.21
UK	4	Seasonal-naive	168	23.28
UK	40	N-BEATS	96	15.58
UK	40	LEAR	319	16.02
UK	40	Seasonal-naive	168	23.22

Table 2.24: Conventional baseline test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; 2025 was not seen during tuning. **Bold** marks the lower error of LEAR and N-BEATS within each scenario-year column. N-BEATS leads on SP and UK in both years, and the seasonal-naive floor trails throughout.

	GER-4		GER-40		SP-4		SP-40		UK-4		UK-40	
Model	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025
LEAR	11.15	**12.90**	**21.56**	**23.57**	46.46	21.68	62.28	33.73	10.21	11.86	16.02	18.35
N-BEATS	**11.09**	13.45	22.20	25.68	**39.11**	**18.46**	**52.15**	**24.94**	**9.06**	**10.49**	**15.58**	**18.28**
Seasonal-naive	30.86	31.77	30.88	31.76	74.76	34.40	74.77	34.40	23.28	25.59	23.22	25.57

of within-horizon price-move directions predicted with the correct sign. Table 2.26 reports these diagnostics at the short horizon ($n = 4$) for the best foundation model (Chronos 2), the best LSTM, and the seasonal-naive floor. Chronos 2 attained the lowest MAE in every regime on all three markets, and its advantage over the best LSTM was largest in the demanding spike and negative-price regimes. Directional accuracy was highest on the smoother GER and UK (near 0.89 and 0.81) and lower on the volatile SP (near 0.67), reflecting the difficulty of anticipating the sign of SP price moves. The same ordering held at $n = 40$ and on the 2025 holdout.

Table 2.25: Membership of the conventional baselines in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Only N-BEATS entered, and only in two scenarios.

Model	GER-4	GER-40	SP-4	SP-40	UK-4	UK-40
LEAR	No	No	No	No	No	No
N-BEATS	No	No	Yes	No	No	Yes
Seasonal-naive	No	No	No	No	No	No

Table 2.26: Regime-segmented test MAE and directional accuracy on the 2024 test year at the short horizon ($n = 4$). Spike: price above the annual mean plus two standard deviations; negative: price below zero; normal: the remainder. Directional accuracy (Dir. acc.) is the fraction of within-horizon move directions predicted with the correct sign. **Bold** marks the best value per column within a market; the SP negative-price column is left unmarked because that regime, with only 12 points, is not interpretable.

Market	Model	MAE by regime			Dir.
		Normal	Spike	Negative	acc.
GER	Chronos 2	**8.15**	**52.38**	**11.17**	**0.89**
	Best LSTM	12.62	73.11	15.29	0.85
	Seasonal-naive	28.61	115.42	39.19	0.86
SP	Chronos 2	**20.41**	**1620.0**	61.5	**0.67**
	Best LSTM	25.84	1838.5	84.8	0.65
	Seasonal-naive	58.25	1858.1	119.4	0.65
UK	Chronos 2	**6.88**	**30.01**	**7.35**	**0.81**
	Best LSTM	11.37	47.87	32.04	0.78
	Seasonal-naive	21.75	76.33	61.52	0.75

Robustness on the 2025 holdout

The 2024 accuracy ranking was largely preserved on the held-out 2025 year. Chronos 2 attained the lowest 2025 MAE in all six scenarios and entered the 2025 MCS [103] in every scenario, being its sole member on GER at both horizons and on UK at $n = 4$, while on SP and on UK at $n = 40$ several foundation models were statistically tied. Crucially, no LSTM configuration and only two conventional baselines (N-BEATS on SP at $n = 4$ and on UK at $n = 40$) entered the 2025 confidence set, as Tables 2.19, 2.25 and 2.28 record, so the foundation-model lead carried over to the unseen year. Market-level holdout drift followed the pattern seen for the LSTM: GER was close to flat for Chronos 2 (from 9.19 to 9.74 at $n = 4$, and from 20.16 to 19.92 at $n = 40$), UK showed small positive drift (from 7.23 to 7.69), and SP showed large negative drift (from 35.03 to 15.56 at $n = 4$, and from 43.03 to 20.66 at $n = 40$), indicating that 2025 was an easier year there. Table 2.27 reports the foundation-model test MAE for both years.

Table 2.27: Foundation-model test MAE on the 2024 test year versus the 2025 holdout year. Models were tuned on 2020–2023 and tested on 2024; the year 2025 was not seen during tuning. **Bold** marks the lowest MAE per column within the family. Chronos 2 remains the most accurate model in almost every column across both years, confirming that its lead is not specific to the 2024 test year.

	GER-4		GER-40		SP-4		SP-40		UK-4		UK-40	
Model	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025	2024	2025
Chronos 2	**9.19**	**9.74**	**20.16**	**19.92**	**35.03**	**15.56**	**43.03**	**20.66**	**7.23**	**7.69**	14.19	**15.88**
Chronos-Bolt (Tiny)	11.00	11.74	20.88	21.52	37.12	16.34	45.73	21.39	8.50	9.29	14.83	16.79
Chronos-Bolt (Small)	9.80	10.80	20.24	20.85	36.28	15.86	44.74	20.99	8.10	8.78	15.04	17.06
Chronos-Bolt (Base)	9.77	10.67	20.59	21.05	35.56	15.87	44.19	20.95	7.93	8.68	14.96	17.36
Moirai-2 (Small)	9.57	10.19	20.68	20.67	36.21	15.83	44.15	21.05	7.66	8.10	**14.13**	16.12
TimesFM 2.5	10.45	11.57	22.27	22.25	36.01	15.78	44.65	21.14	8.17	8.88	15.07	16.58

Table 2.28: Membership of the foundation models in the 2025 out-of-year MCS at the 10% significance level. Yes marks a model inside the 2025 confidence set, No a model excluded as significantly less accurate than the leading model. Chronos 2 entered in every scenario and was the sole member on GER at both horizons and on UK at $n = 4$.

Model	GER-4	GER-40	SP-4	SP-40	UK-4	UK-40
Chronos 2	Yes	Yes	Yes	Yes	Yes	Yes
Chronos-Bolt (Tiny)	No	No	Yes	No	No	Yes
Chronos-Bolt (Small)	No	No	Yes	Yes	No	No
Chronos-Bolt (Base)	No	No	Yes	Yes	No	No
Moirai-2 (Small)	No	No	Yes	Yes	No	Yes
TimesFM 2.5	No	No	Yes	Yes	No	Yes

2.4.5 Discussion

Pre-trained Transformer-based foundation models, in particular Chronos 2, Chronos-Bolt, and Moirai-2, represent a viable and often superior alternative to task-specific recurrent models for day-ahead EPF. Their consistent dominance across all six prediction scenarios, over both the optimised LSTM and the conventional LEAR, N-BEATS, and seasonal-naive baselines and despite operating in a zero-shot regime, suggests that the temporal patterns in electricity price series share sufficient structural similarity with the heterogeneous time-series corpora used during pre-training to enable effective transfer. The regime-segmented analysis reinforced this conclusion, as Chronos 2 retained the lowest error even in the spike and negative-price regimes that are the most demanding for price forecasting. The cross-model performance hierarchy was stable across markets and horizons, indicating that model selection findings are likely to generalise to other day-ahead markets. The compression of performance differences among the top models at longer horizons further suggests that context integration matters more than fine-grained architectural differences when the prediction window widens.

The systematic optimisation of the context window length is essential for the deployment of foundation models. The optimised context lengths varied substantially across models and scenarios, and the market-dependent pattern (optimal context lengths of a broadly similar order in calendar time, roughly one week to a few months, once the sampling interval is accounted for) provides practical guidance: when deploying foundation models on new markets, initial context-length search ranges should be calibrated in calendar time rather than in number of timesteps.

Despite the dominance of foundation models in absolute accuracy, the LSTM experiments from Section 2.3 remain informative. They establish that CV and the CDSC can substantially improve the out-of-sample performance of recurrent models; the within-family spread between the best and worst LSTM configurations reached 5.7% on GER at $n = 4$ and 16.7% on SP at $n = 40$. The systematic comparison of stopping criteria demonstrated that the CDSC is a viable and often superior alternative to standard early stopping, and that the joint optimisation of stopping-criterion parameters alongside model architecture yields configurations that would not be discovered by treating these as fixed design choices.

The performance gap between the best foundation model and the best LSTM was wider on the hourly GER and UK (26.4%–40.8%) than on the half-hourly SP (18.4%–19.0%). This pattern is consistent with the hypothesis that GER and UK, as interconnected European exchanges with hourly resolution, may be more representative of the types of time series present in foundation model pre-training corpora; SP, with its half-hourly resolution and more intense price dynamics, may present a less familiar distributional pattern, partially narrowing the foundation models' advantage. These results are particularly noteworthy given the asymmetry in computational effort between the two model classes. The foundation models operated in a purely zero-shot regime with no task-specific gradient updates; only the context window length was optimised via grid search. The LSTM models, by contrast, benefited from extensive hyperparameter optimisation across architecture, training, and regularisation parameters using extensive Optuna-based hyperparameter optimisation with $R = 3$ repetitions each. The foundation models' superior performance, achieved at a fraction of the computational cost of LSTM training campaigns, demonstrates the practical value of large-scale pre-training on heterogeneous time-series corpora for EPF.

This comparison is, however, between a univariate supervised model (LSTM) and zero-shot models pre-trained on large heterogeneous corpora. The foundation models bring implicit cross-domain knowledge to the forecasting task that the univariate LSTM, by construction, cannot access. The observed performance gap therefore reflects not only an architectural difference but also a fundamental difference in the information available to each model class.

A pervasive pattern in the results was the large negative MAE drift observed across both foundation models and LSTM configurations, particularly on SP: validation MAE values were substantially higher than test MAE values because the validation period (2020–2023) included the highly volatile COVID-19 and energy crisis years, while the 2024 test year represented a return to more moderate price levels. This drift demonstrates that the model selection protocol is conservative, as models selected for good performance during difficult years tend to generalise well to easier years. The held-out 2025 evaluation directly addresses the single-test-year concern: the model rankings and the foundation-model lead carried over to the unseen year, with Chronos 2 remaining inside the 2025 MCS in every scenario and no LSTM configuration entering it, indicating that the reported improvements are not an artefact of the 2024 test year alone.

Chapter 3

Energy Trading Strategy Optimisation for Battery-Enhanced Power Plants

3.1 Introduction

As established in Section 1.3, the profitability of battery-enhanced renewable power plants depends on the energy trading strategy employed on the day-ahead market (DAM) [6]. Given the price forecasts developed in Chapter 2, this chapter presents the optimisation methods I developed to translate predicted prices into optimal hourly buy/sell decisions, spanning two complementary contributions: an adaptive gradient-descent-based trading method that exploits the structure of the energy trading objective (Thesis II.1), and a differential-evolution-based metaheuristic together with a comprehensive system sizing guide (Thesis II.2). Both new methods operate on the same parametrised mathematical model of DC-coupled hybrid inverter systems, which is presented as shared infrastructure in Section 3.2.

The mathematical modelling of battery-enhanced photovoltaic (PV) energy systems requires capturing the interactions among PV generation, battery energy storage system (BESS) dynamics, inverter power limits, and grid exchange. DC-coupled systems, where the PV array and battery share a common DC bus before a single inverter, offer simplified energy flow paths compared to AC-coupled alternatives [46]. Ghenai and Bettayeb [47] described grid-tied hybrid solar inverters combining on-grid and off-grid properties, which form the basis of most commercially available household-sized systems. The model must account for battery capacity constraints (state-of-charge (SoC) limits), storage efficiency losses (round-trip efficiency), inverter power saturation, and the interaction between PV generation and trading decisions. While individual components are well understood, a parametrised model that captures all these constraints in a form suitable as an objective function for hourly trading optimisation, where the decision variable is the continuous hourly energy exchange ratio rather than a binary on/off signal, has not been established in the literature.

Gradient-based methods are the workhorse of continuous optimisation, but their application to energy trading on DAMs requires careful adaptation. The objective function, daily profit as a function of the hourly trading strategy vector, exhibits a complex landscape shaped by battery constraints, storage losses, and the temporal correlation of electricity prices. Standard gradient descent assumes a smooth, unconstrained objective; the presence of hard constraints (battery capacity bounds, inverter power limits) and the discrete nature of the trading horizon (24 hourly slots) necessitates customised approaches. In the energy systems literature, gradient information has been used primarily in the context of neural network (NN) training for electricity price forecasting (EPF) [67] and load forecasting, rather than for the trading optimisation itself. An adaptive step-size gradient descent method that exploits the structure of the energy trading objective, specifically, the initial monotonicity of the profit function with respect to transaction magnitude, has not been proposed.

Differential evolution (DE), introduced by Storn and Price [68], is a population-based metaheuristic for continuous optimisation that has been widely applied in power systems. DE maintains a population of candidate solutions and improves them through mutation, crossover, and selection operations. Its advantages include a small number of control parameters, good convergence properties, and the ability to escape local optima. In power

systems, DE has been applied to a range of optimisation tasks. Noman and Iba [69] applied DE to economic load dispatch problems. Optimal power flow has been addressed by similar means, and Ramli et al. [70] used a multi-objective self-adaptive variant of DE to size a PV/wind/diesel hybrid microgrid system. However, the direct application of DE to the optimisation of hourly energy trading strategies, where each individual in the population represents a complete 24-hour trading schedule and the fitness function incorporates a parametrised system model with storage constraints, has not been explored. Adapting DE to this domain requires encoding the trading strategy as a real-valued vector of hourly exchange ratios and designing mutation operators that respect the system's physical constraints.

The integration of EPF with trading optimisation falls under the broader umbrella of model predictive control (MPC) for energy systems. Unlike receding-horizon MPC, which re-optimises continuously as new measurements arrive, the methods proposed here solve a single day-ahead optimisation over the gate-closure horizon, matching the once-daily bidding structure of the DAM; a receding-horizon extension that re-plans intraday is left as future work. As discussed in Section 1.3, predictive-control approaches have been applied to several adjacent problems. Parisio et al. [64] formulated microgrid operation optimisation as an MPC problem. Zhuo [65] described a microgrid energy-management strategy that combines profit optimisation with battery-life management via approximate dynamic programming. Chen et al. [66] reported a battery management strategy that couples Long Short-Term Memory (LSTM)-based prediction with deep deterministic policy gradient decision-making for personalised microgrid requirements. A common finding is that the accuracy of the price forecast bounds the quality of the control strategy [67]. What remains unexplored is the systematic evaluation of how prediction error propagates through the trading optimisation and affects profit across a wide range of system configurations. Furthermore, while individual studies optimise for a specific system, no comprehensive sizing guide has been published that maps battery capacity, inverter power, and storage efficiency to achievable profit increases under both oracle (perfect information) and predicted (imperfect information) scenarios.

This chapter addresses these gaps. Section 3.2 presents the parametrised mathematical model of DC-coupled hybrid inverter systems that serves as the shared objective function for both optimisation methods. Section 3.3 presents the Optimised Energy Trading by Gradient Descent (OptEnT-GD) algorithm that I developed, a novel adaptive gradient-descent-based method that exploits the structure of the energy trading problem, together with its comparative evaluation across 42 system configurations and a derived implicit-regularisation analysis. Section 3.4 presents the Optimised Energy Trading by Differential Evolution (OptEnT-DE) algorithm that I developed, a differential-evolution-based metaheuristic, together with its benchmarking against OptEnT-GD and a system sizing guide. The formal thesis statements for Thesis II are gathered in Section 5.2, alongside those of Thesis I and Thesis III.

3.2 Parametrised system model

A widely used DC-coupled hybrid grid-tied solar inverter system with BESS was modelled as a proof-of-concept platform. Such systems, originally designed for power-outage protection or grid isolation, can buy, sell, and store energy following a predefined hourly trading strategy. Their operational constraints are well defined, facilitating the creation of an accurate parametric mathematical model. The structure and the fundamental energy and data paths of the proposed system are illustrated in Figure 3.1.

The energy market was modelled by the Hungarian Power Exchange (HUPX) DAM, which is interconnected with most Central European exchanges. The DAM was chosen over the intraday market due to its higher trading volume, which increases the practical relevance of the evaluation.

The optimisation aims to determine the amount of energy to be sold or bought in

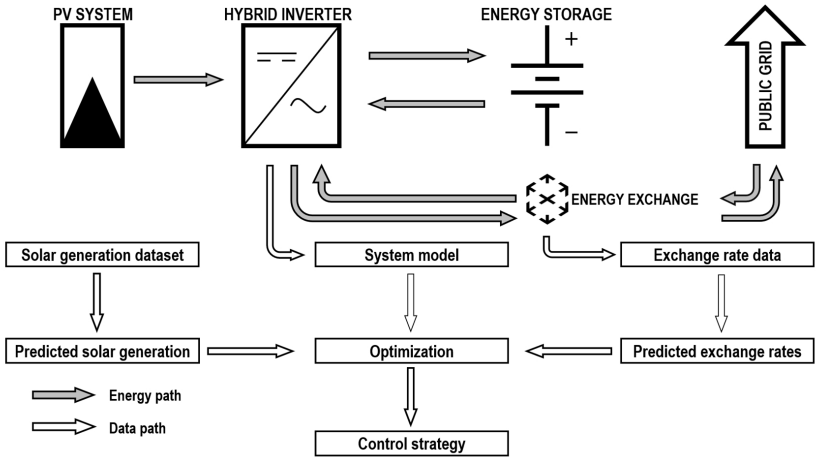


Figure 3.1: Structure of the proposed battery-enhanced PV energy system. The hybrid inverter manages energy flow among the PV array, BESS, the public grid, and the energy exchange. The EPF module provides forecast inputs to the trading strategy optimisation.

every hour over the next n hours (in this work, $n = 24$). The trading ratio in the i-th hour, $x_i \in [-100, 100]$, represents the percentage of the maximum inverter power used for trading: negative values indicate buying energy from the grid, positive values indicate selling. The percentage scale is used throughout the chapter, so that the step sizes $\alpha_{\min}$ of Section 3.3 are expressed in the same units as x_i. The full trading strategy is the vector $\vec{X} = [x_1, x_2, \ldots, x_n]$.

The model deliberately favours simplicity and computational speed over physical fidelity. The trading optimisation framework of Sections 3.3 and 3.4 evaluates the objective function repeatedly for every trading day, so every component of the model has been pared down to the smallest representation that still captures the dominant economic and physical effects. Three scope-defining choices follow directly from this stance and are consequential enough to be stated up front. First, the plant is treated as a price-taker on a single market: only the HUPX DAM is in scope (the intraday and balancing markets are excluded), bids do not move the published clearing price Z_i, every bid clears in full, and buying and selling occur at the same single hourly price, with no bid-ask spread, no asymmetric retail or feed-in tariffs, and no transaction costs, exchange or membership fees, taxes, balancing-market penalties, distribution charges, or transmission losses. Second, each calendar day is optimised independently with a fixed horizon of $n = 24$ hours: there is no rolling-horizon update, no inter-day carry-over of the state of charge, and no multi-day strategy: every day starts from an empty battery ($Q_1 = 0$) and the end-of-day state is left unconstrained, so any energy still stored at hour n is forfeited. This independence is supported a posteriori by the inter-day analysis of Section 3.4.4, where the optimal state of charge at the day boundary was zero in all 42 configurations. Third, no on-site load is included inside the system boundary, so every kilowatt-hour generated, purchased, or stored is accounted for purely through the trading channel. Together these choices reduce the optimisation problem to a sequence of self-contained 24-dimensional subproblems, which is what makes the gradient-based and metaheuristic searches of Sections 3.3 and 3.4 tractable on commodity hardware.

3.2.1 Energy flow equations

The energy transferred to the exchange without adjustments in the i-th hour is:

$$E_i^{\text{base}} = \frac{x_i}{100} \cdot P_{\max}, \tag{3.1}$$

where $P_{\max}$ is the maximum output power of the inverter and the trading slot is one hour, so the energy transferred in a slot is numerically equal to the average power over it (kWh per hour equals kW). Accounting for the round-trip storage efficiency η_{st} ($0 < \eta_{\text{st}} \leq 1$) and the PV generation E_i^{PV}, the adjusted energy transfer becomes:

$$E_i^{\text{adj}} = \eta_{\text{st}} \cdot E_i^{\text{base}} + E_i^{\text{PV}}. \tag{3.2}$$

The battery charge for the $(i{+}1)$-th hour with storage loss is:

$$Q_{i+1}^{\text{adj}} = \begin{cases} \eta_{\text{st}} \cdot (-E_i^{\text{base}}) + Q_i, & \text{if } -E_i^{\text{base}} > 0, \\ -E_i^{\text{base}} + Q_i, & \text{otherwise.} \end{cases} \tag{3.3}$$

Three reductions concentrate in this hour-by-hour energy balance. The storage and inverter losses are aggregated into a single dimensionless coefficient η_{st}, applied once per active leg of the energy path (charge or discharge) rather than split between the two as a true round-trip efficiency, so a kilowatt-hour cycled through the battery incurs the loss exactly once rather than at both ends; η_{st} is held invariant with respect to state of charge, C-rate, ambient temperature, and battery age, and self-discharge during idle hours is neglected. The hourly PV generation E_i^{PV} enters as a known, deterministic, exogenous input: PV forecast uncertainty is not modelled, and module-temperature derating, maximum-power-point-tracking (MPPT) losses, and shading dynamics are absorbed into the value of E_i^{PV} rather than represented explicitly. Finally, power is held constant over each hour, which matches the gate-closure granularity of the DAM and avoids any sub-hourly dispatch dynamics. These three reductions collapse the per-hour update of equations (3.1) to (3.3) to a handful of multiplications and additions, which is the property that makes the inner loop of the optimiser cheap enough to repeat at scale. Note that Equations (3.1) to (3.3) do not include the physical limitations of the system. To accurately model the system, these constraints must be incorporated.

3.2.2 Physical constraints

The physical constraints of the system are modelled in three stages: overcharge protection, depletion protection, and the inverter power limit. First, overcharge protection is applied. If the battery would be overcharged ($Q_{i+1}^{\text{adj}} > Q_{\max}$):

$$E_i^{\text{prot}} = E_i^{\text{adj}} + \frac{Q_{i+1}^{\text{adj}} - Q_{\max}}{\eta_{\text{st}}}, \qquad Q_{i+1}^{\text{prot}} = Q_{\max}. \tag{3.4}$$

Second, depletion protection is applied. If the battery would be depleted ($Q_{i+1}^{\text{adj}} < 0$):

$$E_i^{\text{prot}} = E_i^{\text{adj}} + \eta_{\text{st}} \cdot Q_{i+1}^{\text{adj}}, \qquad Q_{i+1}^{\text{prot}} = 0. \tag{3.5}$$

Third, the inverter power limit is enforced. If the energy to be injected exceeds $P_{\max}$:

$$E_i = \begin{cases} P_{\max}, & \text{if } E_i^{\text{prot}} > P_{\max}, \\ E_i^{\text{prot}}, & \text{otherwise,} \end{cases} \qquad Q_{i+1} = \begin{cases} Q_{i+1}^{\text{prot}} + \dfrac{E_i^{\text{prot}} - P_{\max}}{\eta_{\text{st}}}, & \text{if } E_i^{\text{prot}} > P_{\max}, \\ Q_{i+1}^{\text{prot}}, & \text{otherwise.} \end{cases} \tag{3.6}$$

The constraint-handling logic of equations (3.4) to (3.6) is itself a deliberate simplification. Violations are resolved by a single-pass sequential projection: the battery overcharge and depletion protections are applied first and second, the inverter-saturation correction third, with excess battery flow rerouted to or from the grid and excess inverter flow held in the battery; no iterative solver is invoked to enforce both limits jointly, and after the inverter correction the battery state is clipped to $[0, Q_{\max}]$, so the final state of charge respects the capacity bound even when energy is held back in the battery. The inverter is treated as an ideal symmetric saturator at $\pm P_{\max}$, with no ramp-rate limit, no dead-band, no switching transient, and no power-dependent conversion-efficiency curve. The battery is reduced to an energy reservoir with hard bounds $0 \leq Q_i \leq Q_{\max}$, with calendar ageing, capacity fade, minimum-power thresholds, auxiliary battery-management consumption, and command-tracking error all neglected, and with $Q_{\max}$ held constant over the simulation horizon. No reactive-power, voltage, frequency-support, power-quality, or grid-side network constraints are represented; the point of common coupling is assumed to admit any flow up to $P_{\max}$ in either direction. The single-pass projection is the lever that keeps constraint enforcement linear in the horizon length and side-effect-free, so the projected state at the end of hour i depends only on the projected state at the start, with no fixed-point iteration inside the objective.

3.2.3 Objective function

The daily net market revenue, referred to throughout this chapter as the daily profit under the zero-cost assumptions of Section 3.2, is calculated as:

$$Y = \sum_{i=1}^{n} E_i \cdot Z_i, \tag{3.7}$$

where Z_i is the day-ahead price at hour i. The objective function for both optimisation methods is $F(\vec{X}) = Y$. All profit values throughout this chapter are reported in EUR$*10^{-3}$ (i.e., thousandths of a euro). To prevent overuse of the battery, which is rated for a fixed number of cycles, a penalisation system for battery usage can be introduced; however, this extension was not within the scope of this chapter.

The objective is therefore the single, linear daily revenue of Equation (3.7), with no risk adjustment, no battery-health term, no peak-shaving credit, and no grid-stability incentive. The decision variable $x_i \in [-100, 100]$ is treated as continuous, with no minimum-bid threshold, no integer or binary trading constraint, and no quantisation of the bid quantity. A linear sum over the horizon keeps the per-evaluation cost of $F(\vec{X})$ proportional to n, a property that the finite-difference gradient search of Section 3.3 relies on, since it evaluates the constrained objective repeatedly in every step.

3.3 Adaptive gradient-descent-based energy trading optimisation

3.3.1 Motivation and context

As discussed in the chapter introduction (Section 3.1), gradient-based methods are the dominant tool for continuous optimisation, yet their application to hourly day-ahead energy trading is non-trivial. The objective function (Equation (3.7)) is shaped by hard battery and inverter constraints (Equations (3.4) to (3.6)) and by the temporal correlation of electricity prices, which together produce a non-smooth landscape with multiple local maxima. In the energy systems literature, gradient information has been employed primarily in NN training for EPF [67] and load forecasting, rather than directly in the trading optimisation. Optimisation-based controllers for battery-enhanced systems have instead relied on dynamic programming, Particle Swarm Optimisation (PSO), mixed-integer linear programming, and MPC variants. Levron et al. [53] addressed optimal power flow in microgrids with energy storage via a dynamic programming formulation. Gu et al. [59] applied a mixed-integer linear programming formulation to online dispatch scheduling in combined cooling, heating, and power microgrids. Wu et al. [54] applied PSO to a multi-objective microgrid energy-management dispatch problem that accounts for the power loss of the converters. Baghaee et al. [56] extended the multi-objective PSO formulation to the optimal power management and sizing of a wind/PV microgrid with hydrogen energy storage. Zhuo [65] described a microgrid energy-management strategy that combines profit optimisation with battery-life management via approximate dynamic programming. Parisio et al. [64] formulated microgrid operation optimisation as an MPC problem.

Two domain-specific structural properties of the trading problem remain unexploited by these general-purpose approaches. First, the most profitable transactions take the form of *paired hour modifications*: buy in a low-price hour, store the energy, and sell in a higher-price hour. Modifying only a single hour rarely improves the daily profit because either the battery SoC constraint or the inverter limit is hit; modifying a coupled pair (or a small tuple) of hours in opposite directions is the natural unit of improvement. Second, near a candidate transaction, the daily profit is initially monotonic in the transaction magnitude, with a single local maximum closest to zero before the constraints saturate. An optimiser that searches for the most promising hour-tuple and then exploits this initial monotonicity directly should converge faster, and to a better strategy, than a black-box optimiser that ignores either property. The gap addressed by Thesis II.1 is the absence of a gradient-descent-based trading-optimisation method that explicitly exploits both the multi-hour transaction structure and the initial monotonicity of the profit function.

3.3.2 Method

The OptEnT-GD algorithm is a gradient-descent-based method specifically designed for the energy trading optimisation problem. Its key novelties are: (i) a gradient calculation procedure that evaluates all possible multi-hour transaction combinations to identify the most profitable transaction pair, and (ii) an adaptive step-size calculation that exploits the initial monotonicity of the objective function to find the optimal exploitation of each transaction. The decision variable is the trading strategy vector $\vec{X} = [x_1, x_2, \ldots, x_n]$, where each $x_i \in [-100, 100]$ represents the percentage of inverter power allocated to trading in hour i, on the same scale as the step sizes $\alpha_{\min}$. The objective function is $F(\vec{X}) = Y$ as defined in Equation (3.7).

The OptEnT-GD algorithm proceeds as follows. First, the objective function is evaluated with the initial trading strategy; the search starts from the immediate-dispatch strategy $\vec{X} = \vec{0}$ and terminates when no transaction with positive potential remains, so the iteration cap $i_{\max}^{\mathrm{GD}}$ is a safeguard that was never reached in the experiments. Then, for up to $i_{\max}^{\mathrm{GD}}$

iterations, the gradient calculation procedure (`GradCalc`) identifies the transaction with the highest profit potential by evaluating all possible modifications of γ simultaneously changed hours. The minimum step size $\alpha_{\min}$ serves as the evaluation granularity: transactions whose potential is below this threshold are discarded. If no positive gradient exists, the algorithm terminates and returns the current strategy. Otherwise, the adaptive step-size procedure (`AdaptiveStep`) determines the optimal exploitation of the selected transaction. The complete OptEnT-GD algorithm is summarised in Algorithm 2.

Algorithm 2: OptEnT-GD: Adaptive gradient-descent-based energy trading optimisation. Each iteration selects the most profitable multi-hour transaction and exploits it with an adaptive step size, stopping when no transaction improves the daily profit.

Input: $\vec{X}$ initial trading strategy
$\alpha_{\min}$ minimum step size
n prediction horizon
γ number of simultaneously changed variables
$i_{\max}^{\mathrm{GD}}$ maximum iterations
$F(\vec{X})$ objective function

Output: $\vec{X}_{\mathrm{opt}}$ optimised trading strategy

$\phi_{\mathrm{best}} \leftarrow F(\vec{X})$;
for $i = 1$ **to** $i_{\max}^{\mathrm{GD}}$ **do**
  $\lambda_{\mathrm{gd}} \leftarrow \texttt{GradCalc}(\vec{X}, \alpha_{\min}, n, \gamma, \phi_{\mathrm{best}})$;
  if $\lambda_{\mathrm{gd}} = -1$ **then**
    return $\vec{X}$; `// No improving gradient found`
  $\vec{X} \leftarrow \texttt{AdaptiveStep}(\vec{X}, \alpha_{\min}, n, \gamma, \phi_{\mathrm{best}}, \lambda_{\mathrm{gd}})$;
  $\phi_{\mathrm{best}} \leftarrow F(\vec{X})$;
$\vec{X}_{\mathrm{opt}} \leftarrow \vec{X}$;
return $\vec{X}_{\mathrm{opt}}$

The gradient calculation searches for the most profitable pairs of transactions in which the system buys energy from the grid, stores it in the battery, and later sells it back to the grid. These transactions can be modelled by modifying the control variables for two arbitrary hours in the daily trading strategy ($x_i, x_j \in \vec{X}$, $i \neq j$) in opposite directions. To avoid getting trapped in local maxima, the algorithm was extended to model transactions with more than two elements; the number of variables simultaneously changed is γ ($\gamma \geq 1$).

The `GradCalc` procedure first fills a matrix $\mathbf{X}_\Delta$ with repeating rows of the supplied initial daily trading strategy $\vec{X}$. The matrix is then modified so that each row corresponds to a possible transaction. The key insight is that a single row index i encodes a γ-tuple of hour indices via base-n positional decomposition: for $\gamma = 2$ and $n = 24$, the (1-based) index i encodes two hours as $h_1 = ((i-1) \bmod n) + 1$ and $h_2 = (\lfloor (i-1)/n \rfloor \bmod n) + 1$, yielding n^γ rows covering all possible transaction combinations; rows in which two decoded hour indices coincide encode no transaction and are discarded. The objective function is evaluated with all rows of $\mathbf{X}_\Delta$, and the algorithm returns the index λ_{gd} of the row that resulted in the highest profit increase. If no positive gradients are found, -1 is returned. The complete gradient calculation is given in Algorithm 3. The per-step cost of the gradient calculation is therefore n^γ objective-function evaluations; for the case of $\gamma = 2$ and $n = 24$ used in this work, this is 576 evaluations per step, which is tractable on commodity hardware. The cost grows polynomially with n at fixed γ and exponentially in γ at fixed n, so prediction

horizons substantially longer than 24 hours or transactions with $\gamma > 2$ are out of scope for this formulation without algorithmic refinements such as sampling-based gradient estimation or coordinate-block decomposition.

Algorithm 3: GradCalc: Gradient calculation for multi-hour transactions. All n^γ candidate transactions are evaluated and the index of the one with the largest profit increase is returned, or -1 if none improves the current strategy.

Input: $\vec{X}$ trading strategy
$\alpha_{\min}$ minimum step size
n prediction horizon
γ number of simultaneously changed variables
ϕ_{temp} current best profit
Output: λ_{gd} index of the largest gradient (encoded form)

for $i = 1$ **to** n^γ **do**
  for $j = 1$ **to** n **do**
    $\mathbf{X}_\Delta[i][j] \leftarrow \vec{X}[j]$;
for $i = 1$ **to** n^γ **do**
  $\sigma \leftarrow 1$;
  for $j = 1$ **to** γ **do**
    $h_j \leftarrow (\lfloor (i-1)/n^{j-1} \rfloor \bmod n) + 1$;
    $\mathbf{X}_\Delta[i][h_j] \leftarrow \text{clip}\big(\vec{X}[h_j] + \alpha_{\min} \cdot \sigma,\ -100,\ 100\big)$;
    $\sigma \leftarrow -\sigma$;
for $i = 1$ **to** n^γ **do**
  if *the decoded hour indices* $h_1, \ldots, h_\gamma$ *of row* i *are not distinct* **then**
    $\vec{\phi}_\Delta[i] \leftarrow 0$;
  else
    $\vec{\phi}_\Delta[i] \leftarrow F(\mathbf{X}_\Delta[i][:]) - \phi_{\text{temp}}$;
if $\max(\vec{\phi}_\Delta) \leq 0$ **then**
  return -1;
$\lambda_{\text{gd}} \leftarrow \arg\max(\vec{\phi}_\Delta)$;
return λ_{gd}

A transaction has two properties: potential and exploitability. The potential is the calculated gradient, and the exploitability is the best step size in the best direction. An advantage of the proposed gradient calculation is that it discards transactions with exploitability below $\alpha_{\min}$, offering an adjustable trade-off. By finding the optimal minimum step size, faster convergence can be achieved and more profitable trading strategies identified. Since $\alpha_{\min}$ is a parameter rather than an algorithmically computed value, it can be tuned as a hyperparameter.

When a positive gradient is found, the adaptive step-size procedure (`AdaptiveStep`) determines the optimal exploitation. The heuristic is based on the objective function's initial monotonic increase, which means the maximum profit is the local maximum closest to zero. Starting from zero, the step size is increased with accelerating increments (i.e., $1 \cdot \alpha_{\min}, 3 \cdot \alpha_{\min}, 6 \cdot \alpha_{\min}, 10 \cdot \alpha_{\min}, \ldots$) until the first decrease in profit is detected. The step size is then refined by reducing it in increments of $\alpha_{\min}$ until the local maximum is located. This two-phase approach balances computational cost with precision. The complete adaptive

step-size procedure is given in Algorithm 4.

Algorithm 4: AdaptiveStep: Adaptive step-size calculation. Exploiting the initial monotonicity of the profit function, an accelerating search first overshoots the optimum and a refinement phase then locates the local maximum nearest zero.

Input: $\vec{X}$ trading strategy
$\alpha_{\min}$ minimum step size
n prediction horizon
γ number of simultaneously changed variables
ϕ_{best} current best profit
λ_{gd} index of the best gradient
$F(\vec{X})$ objective function

Output: $\vec{X}_{\text{astep}}$ modified trading strategy

$\phi_{\text{temp}} \leftarrow \phi_{\text{best}}$;
// Phase 1: accelerating search, stopped at the first decrease
for $i = 1$ **to** $\lfloor 100/\alpha_{\min} \rfloor$ **do**
  $\phi_{\text{prev}} \leftarrow \phi_{\text{temp}}$;
  $\vec{X} \leftarrow \texttt{Step}(\lambda_{\text{gd}}, \vec{X}, i \cdot \alpha_{\min}, n, \gamma, -1)$;
  $\phi_{\text{temp}} \leftarrow F(\vec{X})$;
  if $\phi_{\text{prev}} > \phi_{\text{temp}}$ **then**
    break;
// Phase 2: refinement back towards the local maximum
for $i = 1$ **to** 100 **do**
  $\phi_{\text{prev}} \leftarrow \phi_{\text{temp}}$;
  $\vec{X} \leftarrow \texttt{Step}(\lambda_{\text{gd}}, \vec{X}, \alpha_{\min}, n, \gamma, +1)$;
  $\phi_{\text{temp}} \leftarrow F(\vec{X})$;
  if $\phi_{\text{prev}} > \phi_{\text{temp}}$ **then**
    $\vec{X} \leftarrow \texttt{Step}(\lambda_{\text{gd}}, \vec{X}, \alpha_{\min}, n, \gamma, -1)$;
    $\phi_{\text{best}} \leftarrow F(\vec{X})$;
    $\vec{X}_{\text{astep}} \leftarrow \vec{X}$;
    return $\vec{X}_{\text{astep}}$;
// Fallback: refinement reached the bound without a decrease
$\phi_{\text{best}} \leftarrow \phi_{\text{temp}}$;
$\vec{X}_{\text{astep}} \leftarrow \vec{X}$;
return $\vec{X}_{\text{astep}}$

The `Step` procedure takes the encoded gradient index λ_{gd}, the step direction σ, and the step size, then performs a step in the direction of the best gradient. The "decoding" process recovers the indices of the hours that make up the transaction: for each of the γ variables, the hour index is extracted from λ_{gd} via modular arithmetic ($h_j = (\lfloor (\lambda_{\text{gd}} - 1)/n^{j-1} \rfloor \bmod n) + 1$, the same decoding as in `GradCalc`), and the corresponding element of the trading strategy is incremented by $\alpha \cdot \sigma$, with alternating signs (σ flips between $+1$ and -1 after each variable). Because `Step` adds to the current strategy cumulatively, successive calls in the first phase of `AdaptiveStep` with step sizes $1 \cdot \alpha_{\min}, 2 \cdot \alpha_{\min}, 3 \cdot \alpha_{\min}, \ldots$ produce total displacements of $1 \cdot \alpha_{\min}, 3 \cdot \alpha_{\min}, 6 \cdot \alpha_{\min}, 10 \cdot \alpha_{\min}, \ldots$, yielding the accelerating search described above.

3.3.3 Experimental methodology

The system model was evaluated across 42 configurations comprising all permutations of the parameters listed in Table 3.1. The constraint PV = $P_{\max}$ was applied so that the peak PV output matches the inverter capacity. This is an experimental-campaign convention rather than a model assumption: the equations of Sections 3.2.1 and 3.2.2 accept any $E_i^{\mathrm{PV}} \leq P_{\max}$, and decoupling the PV peak from the inverter rating only requires changing the input series. One year of historical HUPX data, divided into 24-hour chunks, served as input, and daily results were averaged over the year.

Table 3.1: Variables and their values evaluated across all scenarios. All possible combinations of system parameters were tested with each trading strategy and optimisation method. The three inverter powers, seven battery capacities, and two storage efficiencies yield the 42 system configurations used throughout the chapter.

Variable	Value
System parameters	
Maximum inverter power ($P_{\max}$)	3 kW, 4 kW, 5 kW
Battery capacity ($Q_{\max}$)	3 kWh, 4.5 kWh, 6 kWh, 7.5 kWh, 9 kWh, 10.5 kWh, 12 kWh
Storage efficiency (η_{st})	0.95, 1.0
Trading strategy	Immediate-dispatch (energy sold immediately), Oracle (oracle day-ahead prices), Predicted (LSTM-forecasted prices)
Optimisation method	OptEnT-GD, OptEnT-DE

Three scenarios were tested: (i) *immediate-dispatch*, where generated energy is sold immediately (baseline); (ii) *oracle*, where the trading strategy is optimised using actual day-ahead prices known in hindsight; and (iii) *predicted*, where the LSTM-predicted prices from Section 2.3 are used. The oracle scenario does not reflect real-world operations, since day-ahead prices for the next day are unknown and must be predicted; however, it serves to assess the algorithms' potential and to provide a basis for quantifying the effect of prediction error.

For OptEnT-GD, the minimum step size $\alpha_{\min}$ was treated as a tunable hyperparameter, with six values evaluated across all 42 configurations: $\alpha_{\min} \in \{0.1, 1, 5, 10, 25, 50\}$. The effect of $\alpha_{\min}$ was analysed in two ways: (i) the number of configurations for which a given step size achieved the best profit, and (ii) the pairwise relative profit differences between step sizes, averaged across configurations.

3.3.4 Results

In the oracle scenario, the results are presented in Table 3.2. The minimum step size $\alpha_{\min} = 5$ achieved the best profit with the largest number of configurations (14), while $\alpha_{\min} = 25$ achieved the best profit with 13 configurations and yielded the highest average profit across the 42 configurations.

In the predicted scenario, the results are presented in Table 3.3. The optimal step size shifted to $\alpha_{\min} = 50$, which achieved the best profit with 22 of the 42 configurations and the highest average profit.

The increase in profit relative to the immediate-dispatch baseline was analysed across all system configurations. In the oracle scenario ($\alpha_{\min} = 25$, $\eta_{\mathrm{st}} = 0.95$), the smallest daily profit increase over the immediate-dispatch baseline was 141.78 EUR$*10^{-3}$ (system with $Q_{\max} = 3$ kWh, $P_{\max} = 5$ kW), while the largest was 465.48 EUR$*10^{-3}$ (system with

Table 3.2: Results of OptEnT-GD in the oracle scenario, grouped by minimum step size. *No. of best results* indicates the number of configurations where the step size achieved the highest profit. Cross-tabulation values show the average relative profit difference between step sizes compared to the immediate-dispatch baseline. The step size $\alpha_{\min} = 5$ won the most configurations (14), whereas $\alpha_{\min} = 25$ gave the highest average profit across the 42 configurations.

Oracle	$\alpha_{\min}$=0.1	$\alpha_{\min}$=1	$\alpha_{\min}$=5	$\alpha_{\min}$=10	$\alpha_{\min}$=25	$\alpha_{\min}$=50
No. of best results	0	3	14	7	13	7
$\alpha_{\min} = 1$	0.10%	–	–	–	–	–
$\alpha_{\min} = 5$	0.41%	0.31%	–	–	–	–
$\alpha_{\min} = 10$	0.89%	0.78%	0.47%	–	–	–
$\alpha_{\min} = 25$	1.59%	1.49%	1.18%	0.71%	–	–
$\alpha_{\min} = 50$	1.49%	1.39%	1.08%	0.61%	−0.10%	–

Table 3.3: Results of OptEnT-GD in the predicted scenario, grouped by minimum step size. The optimal step size shifted from $\alpha_{\min} = 25$ in the oracle scenario to $\alpha_{\min} = 50$ in the predicted scenario.

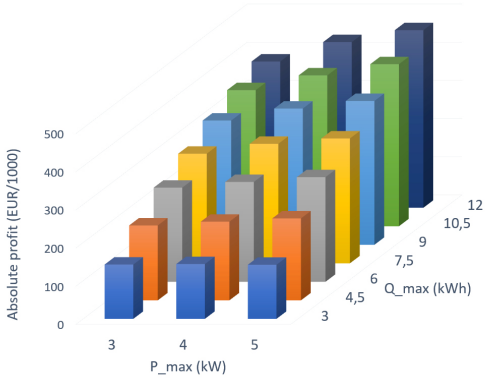
Predicted	$\alpha_{\min}$=1	$\alpha_{\min}$=5	$\alpha_{\min}$=10	$\alpha_{\min}$=25	$\alpha_{\min}$=50
No. of best results	1	1	9	12	22
$\alpha_{\min} = 5$	0.01%	–	–	–	–
$\alpha_{\min} = 10$	0.03%	0.02%	–	–	–
$\alpha_{\min} = 25$	0.21%	0.20%	0.17%	–	–
$\alpha_{\min} = 50$	0.74%	0.73%	0.71%	0.54%	–

$P_{\max} = 5$ kW, $Q_{\max} = 12$ kWh). The absolute and percentage profit increases are shown in Figure 3.2.

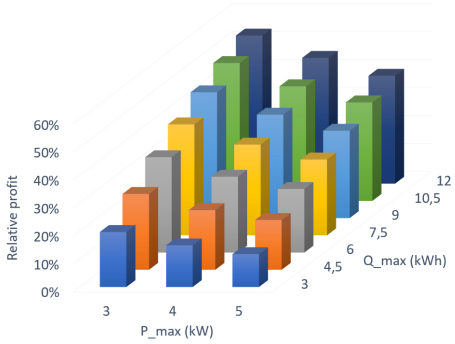
The lowest percentage profit was 19.66% (system with $P_{\max} = 5$ kW, $Q_{\max} = 3$ kWh), and the largest was 52.94% (system with $P_{\max} = 3$ kW, $Q_{\max} = 12$ kWh), exceeding 50% over the immediate-dispatch case. The percentage profit increased monotonically with the battery-capacity-to-inverter-power ratio across the tested configurations.

In the predicted scenario ($\alpha_{\min} = 50$), the system behaved similarly to the oracle case (Figure 3.3). The difference between two systems with the same battery size ($Q_{\max} = 12$ kWh) but different inverter power ($P_{\max} = 3$ kW vs. $P_{\max} = 5$ kW) was 82.51 EUR$*10^{-3}$ in the oracle case and 61.98 EUR$*10^{-3}$ in the predicted case. With predicted day-ahead prices, the OptEnT-GD method achieved up to 46.26% profit increase compared to the immediate-dispatch baseline (system with $P_{\max} = 3$ kW, $Q_{\max} = 12$ kWh).

The percentage profit difference between the predicted and the oracle scenarios is shown in Figure 3.4. Across the 42 configurations, the prediction error lowered the percentage profit increase over the immediate-dispatch baseline by 1.93 to 6.93 percentage points relative to the oracle scenario. In the configuration with the largest difference ($P_{\max} = 3$ kW, $Q_{\max} = 12$ kWh), the profit increase with EPF was still 46.26%.

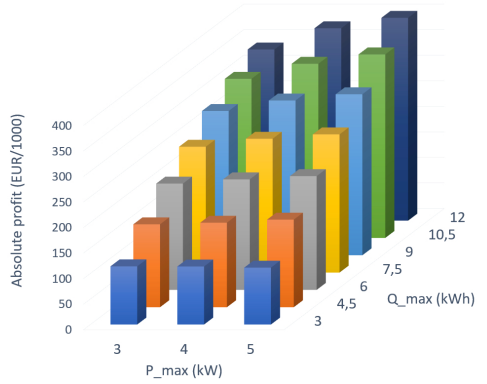


(a) Daily profit increase over immediate dispatch (EUR$*10^{-3}$) by battery capacity and inverter power

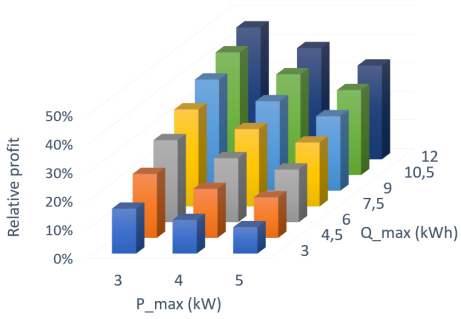


(b) Percentage profit increase relative to immediate-dispatch baseline

Figure 3.2: Absolute (a) and percentage (b) profit increase of OptEnT-GD over the immediate-dispatch baseline in the oracle scenario ($\eta_{st} = 0.95$, PV = P_{max}, $\alpha_{min} = 25$). The percentage gain rises monotonically with the battery-capacity-to-inverter-power ratio, exceeding 50% for the largest battery with the smallest inverter.



(a) Daily profit increase over immediate dispatch (EUR$*10^{-3}$) by battery capacity and inverter power



(b) Percentage profit increase relative to immediate-dispatch baseline

Figure 3.3: Absolute (a) and percentage (b) profit of OptEnT-GD compared to the immediate-dispatch baseline in the predicted scenario ($\eta_{st} = 0.95$, PV = P_{max}, $\alpha_{min} = 50$). With LSTM-predicted prices the pattern of the oracle case is preserved, reaching up to 46.26% profit increase over the immediate-dispatch baseline.

3.3.5 Discussion

The results demonstrate that OptEnT-GD achieves substantial profit increases across all 42 tested system configurations, with the largest gains observed for systems with high battery-capacity-to-inverter-power ratios. The percentage profit increase exceeded 50% for the most favourable configuration, confirming that the multi-hour transaction structure and the adaptive step-size function together exploit the available arbitrage opportunities effectively. This monotonic dependence on the battery-capacity-to-inverter-power ratio is consistent with the structure of the objective function: a larger battery enables more energy to be shifted between low- and high-price hours within the inverter power limit,

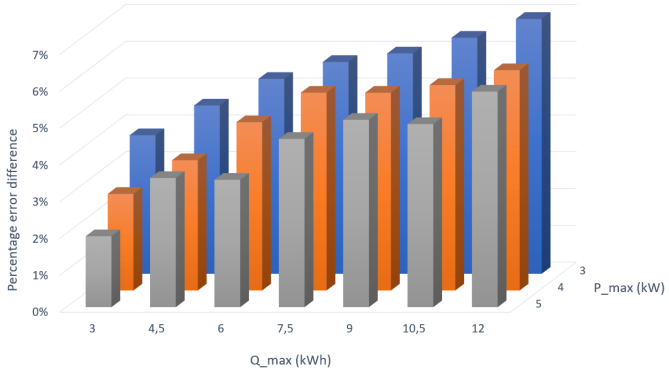


Figure 3.4: Percentage profit difference between the predicted and oracle scenarios for different system configurations (PV = $P_{\max}$). Across the 42 configurations, the prediction error lowers the profit increase over the immediate-dispatch baseline by only 1.93 to 6.93 percentage points.

while smaller batteries saturate the storage constraint sooner and forfeit later transactions.

The shift in optimal minimum step size from $\alpha_{\min} = 25$ in the oracle scenario to $\alpha_{\min} = 50$ in the predicted scenario is consistent with an implicit regularisation effect of $\alpha_{\min}$. Larger step sizes force the algorithm to discard fine-grained transactions whose potential is below the chosen threshold; under prediction noise, such fine-grained transactions are also the ones most sensitive to small price errors, so discarding them produces more general trading strategies that transfer better from forecasted to realised prices. This makes $\alpha_{\min}$ a tunable knob trading transaction granularity for robustness to forecast error, and provides a practical guideline for deployment: the optimal $\alpha_{\min}$ should be re-tuned whenever the upstream price-forecasting pipeline changes.

The modest prediction-error penalty (a 1.93 to 6.93 percentage-point reduction of the profit increase) indicates that the combined LSTM forecasting (Section 2.3) and OptEnT-GD optimisation pipeline is robust to prediction error within the tested simulation setting: even with imperfect price forecasts, the headline 46.26% profit increase over immediate dispatch is preserved in the most favourable configuration. The link between the battery-capacity-to-inverter-power ratio and profit gain identified here motivates the comprehensive sizing analysis carried out in Section 3.4, where it is generalised into a system sizing guide.

A limitation of the evaluation is that the benchmark baseline is the no-trading immediate-dispatch strategy. Comparisons against published optimisation-based trading methods (mixed-integer linear programming, MPC variants, dynamic programming, PSO) on the same parametrised system model were not part of this work, and are reserved for future work alongside the LSTM-forecast-coupled deployment study.

3.4 Differential-evolution-based energy trading optimisation and system sizing guide

3.4.1 Motivation and context

The OptEnT-GD algorithm of Section 3.3 relies on two structural assumptions: that the daily profit function is initially monotonic in the magnitude of each candidate transaction, and that the most profitable improvements take the form of paired hour modifications enumerable by exhaustive multi-hour gradient evaluation. Both assumptions are well supported under the parametrised model of Section 3.2, but they constrain the applicability of the method to objective functions that share these properties. Modified objectives that incorporate, for example, battery-degradation penalties, time-varying inverter efficiency, or non-symmetric buy/sell tariffs may break the initial-monotonicity property and render the gradient-based search suboptimal. A complementary, derivative-free optimiser that does not depend on either assumption is therefore needed to widen the class of objective functions over which the trading framework remains usable.

DE, introduced by Storn and Price [68], is a population-based metaheuristic for continuous optimisation that has been widely applied in power systems. Noman and Iba [69] applied DE to economic load dispatch problems, while optimal power flow has also been addressed by DE-based methods. Ramli et al. [70] used a multi-objective self-adaptive variant of DE to size a PV/wind/diesel hybrid microgrid system. DE explores the search space through mutation, crossover, and selection operations, requires only function evaluations (no derivatives), and is a population-based global-search metaheuristic able to escape local optima with a small number of control parameters. Its direct application to the optimisation of hourly trading strategies, where each individual encodes a complete 24-hour schedule and the fitness function is the constrained system model of Section 3.2, has not been reported.

A second open question concerns the relationship between system parameters and achievable profit. The 42-configuration evaluation of Section 3.3.4 reveals a clear monotonic link between the battery-capacity-to-inverter-power ratio and the relative profit gain, but the literature lacks a comprehensive sizing guide that consolidates this relationship across both optimisation methods, the oracle and predicted scenarios, and the question of whether inter-day energy storage offers further profit. The two gaps addressed by Thesis II.2 are therefore: (i) the absence of a DE-based trading optimiser for battery-enhanced PV energy systems, and (ii) the absence of a system sizing guide that unifies the empirical findings of both OptEnT-GD and OptEnT-DE.

3.4.2 Method

The OptEnT-DE algorithm adapts the general-purpose DE metaheuristic [68] to the energy trading domain. DE is a population-based global optimiser that uses mutation, crossover, and selection to explore potentially nonlinear, non-differentiable solution spaces. The evolutionary operators described below are the standard DE operators; the contribution is the problem-specific formulation, in which each individual encodes a complete 24-hour trading schedule and fitness is the constrained profit objective of Section 3.2, identical to the one optimised by OptEnT-GD so that the two methods can be compared on equal terms.

The population $\mathcal{P}$ consists of s individuals, where each individual $\vec{X}_j = [x_1, x_2, \ldots, x_n]$ represents a candidate trading strategy. The OptEnT-DE algorithm first randomly initialises the population, with each element drawn uniformly from $[-100, 100]$, and then evaluates the objective function (Equation (3.7)) for all individuals. Then, for each generation and each $\vec{X}_j$, the mutation, crossover, and selection operations are performed.

Mutation: Three distinct candidates $\vec{X}_a$, $\vec{X}_b$, and $\vec{X}_c$ are selected from the population. A mutant individual $\vec{m}$ is created as:

$$\vec{m} = \vec{X}_a + F_{\text{de}} \cdot (\vec{X}_b - \vec{X}_c), \tag{3.8}$$

where F_{de} is the scale (mutation) factor; the mutant is clipped element-wise to the feasible range $[-100, 100]$ before crossover.

Crossover: A trial individual $\vec{u}$ is constructed via binomial crossover with rate CR. First, a random index r is picked from the dimensionality of the problem ($1 \leq r \leq n$). Then, for each element k:

$$\vec{u}[k] = \begin{cases} \vec{m}[k], & \text{if } \text{rand}(0,1) < CR \text{ or } k = r, \\ \vec{X}_j[k], & \text{otherwise.} \end{cases} \tag{3.9}$$

The mandatory index r ensures that the trial individual differs from the parent in at least one element.

Selection: The trial replaces the parent if and only if $F(\vec{u}) > F(\vec{X}_j)$ (greedy selection for maximisation). After all generations are complete, the individual with the highest fitness in the final population is returned as the optimised trading strategy. The complete OptEnT-DE algorithm is given in Algorithm 5.

Algorithm 5: OptEnT-DE: Differential-evolution-based energy trading optimisation. A population of 24-hour trading schedules is evolved through mutation, crossover, and greedy selection, and the fittest schedule in the final population is returned.

Input: n prediction horizon
s population size
$i_{\max}^{\mathrm{DE}}$ maximum iterations
CR crossover rate
F_{de} mutation factor
$F(\vec{X})$ objective function

Output: $\vec{X}_{\mathrm{opt}}$ optimised trading strategy

`// Initialise population`

for $i = 1$ **to** s **do**
  for $j = 1$ **to** n **do**
    $\vec{X}_i[j] \leftarrow \mathrm{rand}([-100, 100])$;
  $\phi_i \leftarrow F(\vec{X}_i)$;

`// Evolution`

for $g = 1$ **to** $i_{\max}^{\mathrm{DE}}$ **do**
  for $j = 1$ **to** s **do**
    Pick three random distinct candidates $\vec{X}_a, \vec{X}_b, \vec{X}_c$ from $\mathcal{P}$;
    $\vec{m} \leftarrow \mathrm{clip}\big(\vec{X}_a + F_{\mathrm{de}} \cdot (\vec{X}_b - \vec{X}_c),\ -100,\ 100\big)$;
    $k_{\mathrm{rand}} \leftarrow \mathrm{randint}(1, n)$;
    for $k = 1$ **to** n **do**
      if $\mathrm{rand}(0, 1) < CR$ ***or*** $k = k_{\mathrm{rand}}$ **then**
        $\vec{u}[k] \leftarrow \vec{m}[k]$;
      else
        $\vec{u}[k] \leftarrow \vec{X}_j[k]$;
    if $F(\vec{u}) > \phi_j$ **then**
      $\vec{X}_j \leftarrow \vec{u}$;
      $\phi_j \leftarrow F(\vec{u})$;
$\vec{X}_{\mathrm{opt}} \leftarrow \vec{X}_{\arg\max(\vec{\phi})}$;
return $\vec{X}_{\mathrm{opt}}$

3.4.3 Experimental methodology

The OptEnT-DE algorithm was evaluated across the same 42 system configurations as OptEnT-GD (Table 3.1), with predicted day-ahead prices supplied by the LSTM model of Section 2.3. The DE control parameters were set following the standard recommendations of Storn and Price [68]: scale factor $F_{\text{de}} = 0.8$ and crossover rate $CR = 0.9$. The population size $s = 50$ was chosen as a practical compromise between the recommended $s = 10 \cdot n = 240$ and the per-day evaluation budget, with 100 stochastic restarts compensating for the reduced population. Due to the stochastic nature of DE, the algorithm was run 100 times for each day, and the resulting daily profits were summarised by their statistical properties (minimum, mean, maximum) before being averaged over the year.

To compare OptEnT-DE against OptEnT-GD under the same predicted-price conditions, the OptEnT-GD runs from Section 3.3 (using the optimal $\alpha_{\text{min}} = 50$ from Table 3.3) were used as the reference. The two methods were compared in absolute profit on the same daily price series.

To investigate whether inter-day energy storage offers an additional profit, the prediction and optimisation window was extended to 48 hours, and the optimal SoC after the first 24 hours was examined across all 42 configurations.

3.4.4 Results

The OptEnT-DE algorithm achieved approximately 3% lower maximum profit than OptEnT-GD, most pronounced in the large-battery configurations. The spread between the minimum and maximum profit across the 100 stochastic runs averaged 25 to 30%. Detailed results are presented for the most informative case: the largest battery configuration ($Q_{\text{max}} = 12$ kWh, $\eta_{\text{st}} = 0.95$) with predicted day-ahead prices, where the difference in absolute profit is most pronounced (Table 3.4).

Table 3.4: Comparison of OptEnT-DE (100 runs) and OptEnT-GD with predicted day-ahead prices, $Q_{\text{max}} = 12$ kWh, $\eta_{\text{st}} = 0.95$. OptEnT-GD yields about 3% higher profit than the best OptEnT-DE run in this large-battery case.

	$P_{\text{max}} = 3$ kW	$P_{\text{max}} = 4$ kW	$P_{\text{max}} = 5$ kW
OptEnT-DE min	726.25	968.33	1210.41
OptEnT-DE mean	756.37	1001.64	1244.94
OptEnT-DE max	1027.48	1301.46	1555.69
OptEnT-GD	1059.27	1342.29	1604.08

The 42-configuration evaluation produced the following relationships between system parameters and profit:

- Systems with larger battery-capacity-to-inverter-power ratios ($Q_{\text{max}}/P_{\text{max}}$) yielded proportionally greater percentage profit increases over the immediate-dispatch baseline. The maximum percentage profit reached 52.94% (oracle scenario, $P_{\text{max}} = 3$ kW, $Q_{\text{max}} = 12$ kWh, $\eta_{\text{st}} = 0.95$) and 46.26% under predicted prices for the same configuration.
- Across the two optimisation methods, OptEnT-GD produced higher maximum profit than OptEnT-DE, by approximately 3% in the large-battery configurations.
- Across the oracle and predicted scenarios, the prediction error lowered the percentage profit increase by 1.93 to 6.93 percentage points.

For the inter-day storage analysis, the optimal SoC after the first 24 hours of the 48-hour optimisation window was SoC = 0 for every one of the 42 configurations.

3.4.5 Discussion

Although OptEnT-DE underperformed OptEnT-GD in maximum profit by approximately 3%, the gap is small enough that OptEnT-DE remains a viable alternative when gradient information is unreliable or when the objective function violates the initial-monotonicity assumption that OptEnT-GD exploits. The high variance (25 to 30% spread between minimum and maximum profit across 100 runs) reflects the stochastic nature of the algorithm and suggests that a larger population size, a longer evolution budget, or an adaptive parameter schedule could narrow this spread. In deployment terms, OptEnT-GD exploits the structure of the objective to converge in a small, deterministic number of evaluations per day, whereas OptEnT-DE requires a population evaluated over many generations and several independent runs to control its variance; both are nonetheless executed once per day rather than in real time, so neither imposes a latency constraint on industrial deployment. The ability of DE to explore non-convex and non-differentiable regions of the search space makes it the natural fallback method for variants of the trading problem that introduce battery-degradation penalties, time-varying inverter efficiency, or other modifications that break the assumptions exploited by OptEnT-GD.

The system sizing analysis condenses the empirical findings of both optimisers into a practical design guide. The headline result, that the percentage profit increase scales monotonically with the battery-capacity-to-inverter-power ratio and reaches more than 50% in the oracle scenario and 46.26% under predicted prices, identifies this ratio as the strongest empirical predictor of the achievable trading uplift among the parameters varied. Combined with the consistent 1.93 to 6.93 percentage-point prediction-error penalty, this gives system designers a quantitative criterion for sizing battery-enhanced PV installations: the marginal profit gain from increasing $Q_{\max}$ at a fixed $P_{\max}$ can inform a subsequent economic sizing analysis once investment costs, degradation, and discounting are supplied. For commercially available system parameters in the tested range, the OptEnT-GD method with $\alpha_{\min} = 50$ and the LSTM-based EPF from Section 2.3 provides the best trade-off between profit maximisation and robustness to prediction error; OptEnT-DE is the recommended fallback when the objective function deviates from the parametrised form of Section 3.2.

The finding that the optimal SoC after the first 24 hours of a 48-hour optimisation window was zero in every one of the 42 configurations indicates that, under the tested HUPX market conditions and the parametrised system model of Section 3.2, inter-day energy storage offers no additional profit beyond what is captured by the 24-hour optimisation horizon. This justifies the choice of $n = 24$ for both methods and bounds the depth of forward planning required by the trading optimiser. The conclusion is conditional on the model: introducing battery-degradation penalties, asymmetric tariffs, or longer-horizon weather correlations could change the picture, and motivates extending the framework with such terms as future work.

The complete trading framework, comprising the parametrised system model of Section 3.2, the gradient-based optimiser OptEnT-GD (Section 3.3), the differential-evolution optimiser OptEnT-DE, and the sizing guide derived above, is modular and customisable, and provides a basis that could be adapted and validated for household-sized PV power plants, virtual power plants, and energy communities.

Chapter 4

Physical Design Optimisation of Photovoltaic Energy Systems

4.1 Introduction

As established in Section 1.4, the physical design of photovoltaic (PV) energy systems under real-world operational, regulatory, and sustainability constraints determines the energy yield, economic return, and environmental impact of renewable installations. Liu et al. [71] report a techno-economic simulation and optimisation of a residential grid-connected PV energy system for the Queensland climate, taking both economic metrics and carbon dioxide emissions into account. Fan and Xia [72] present a multi-objective optimisation model for building-integrated PV retrofits, evaluated on a large apartment complex. While Chapter 2 addressed the prediction of day-ahead electricity prices and Chapter 3 addressed the optimisation of trading strategies for battery-enhanced power plants, this chapter targets the upstream problem: designing the PV energy system itself so that it is sized, configured, and oriented for its specific operational context. The chapter spans two complementary contributions: a consumption-aware optimisation algorithm for industrial multi-string PV energy systems under feed-in restrictions (Thesis III.1), and a comprehensive energy system design methodology for semi-autonomous structures equipped with PV generation and a battery energy storage system (BESS) (Thesis III.2).

Traditional PV energy system design relies on matching a system's annual sum of output to the site's annual sum of consumption (the Sum Yearly Consumption Matching (SYCM) method) [71]. This approach ignores the temporal mismatch between generation and consumption, which is particularly problematic for industrial prosumers prohibited from injecting excess energy into the public grid: surplus generation is curtailed by feedback protection rather than sold, making the matching approach economically wasteful [7]. An alternative approach limits the energy lost to feedback protection (E_{lost} limit method) but tends to produce undersized systems that do not significantly reduce grid reliance [7]. Both approaches lack a structured optimisation framework that simultaneously balances system size, economic return, and sustainability objectives.

A more refined approach compares simulated PV output against measured consumption data at high temporal resolution (typically 15-minute intervals) to compute energy-balance-aware metrics. The industry-standard Photovoltaic Geographical Information System (PVGIS) tool [110] simulates PV output for arbitrary surface orientations and panel types from satellite-derived irradiance data. For economic evaluation, decision-makers commonly use net present value (NPV), internal rate of return (IRR), and levelised cost of electricity (LCOE). Liu et al. [71] apply techno-economic metrics together with carbon dioxide emissions in their analysis of residential grid-connected PV energy systems. Diaf et al. [73] use LCOE as part of a techno-economic optimisation of a hybrid PV/wind system on Corsica Island. Fan and Xia [72] adopt NPV and the payback period as the economic metrics in their multi-objective PV retrofitting model. Sustainability metrics, particularly the CO_2 emissions associated with grid consumption, are gaining importance as corporate environmental targets become binding. The hourly specific CO_2 intensity of grid electricity can be computed by combining the energy mix of a country with per-technology life-cycle emission factors. The hourly energy-mix data (the share of generation by production type) is published by the European Network of Transmission System Operators for Electricity

(ENTSO-E) transparency platform [111]. The Intergovernmental Panel on Climate Change (IPCC) AR5 Mitigation report [112] provides technology-specific cost and performance parameters, including the life-cycle emission factors used to translate each production type into a CO_2 intensity. The United Nations Economic Commission for Europe (UNECE) integrated life-cycle assessment of electricity sources [113] offers a more recent set of factors covering the same generation technologies for the UNECE region.

Despite the availability of these data sources and metrics, no optimisation algorithm has been proposed that designs multi-string PV energy systems by optimising an objective function that is unimodal in string direct current (DC) output power under feed-in restrictions, with per-string logarithmic convergence and support for multiple surfaces with different orientations. The dichotomous search algorithm is a classical method for finding the optimum of a unimodal function on a bounded interval [114]. While well established in numerical methods, its application to PV string sizing, where each string's peak power is the decision variable and the objective function couples PV output simulation, consumption data, tariff structure, and economic or sustainability metrics, has not been explored. The key insight is that, for a fixed load profile and tariff, the NPV exhibits unimodality with respect to string power: undersized systems yield insufficient savings, oversized systems waste energy through curtailment, and a single optimum exists between the two regimes.

Semi-autonomous structures (those that operate independently from the public grid for most of the year while retaining a limited grid backup) pose distinct energy system design challenges. Such systems must jointly optimise PV generation, BESS, and thermal management under tight physical and regulatory constraints, including limited grid connection capacity, prohibition of feed-in, weight or space limitations on storage, and the need for high self-reliance ratios. Smart-home and microgrid review articles provide the broader background [80]. Akinyele et al. [81] presented a case study on an autonomous home demonstrating feasibility. Aguilar et al. [115] proposed consumption control methods to better utilise non-controllable renewable generation. Tipantuna and Hesselbach [116] proposed an IoT-enabled adaptive self-powered renewable energy management system.

The floating-home case study considered later in this chapter is a representative semi-autonomous structure for which these constraints are particularly pronounced. Such a structure can be rotated within its mooring and design constraints, enabling optimisation of PV panel orientation; its insulation and sun exposure differ from land-based buildings, affecting the thermal balance; the platform's weight capacity constrains battery size; and limited grid connectivity demands the highest possible self-reliance ratio. Related research has progressed on the structural side. Ishaque et al. [84] described a structural design for a low-cost floating house in which safety was the primary criterion. Cahyono [117] reported a stability test of a floating-house platform with a lightweight concrete cover made of expanded polystyrene (EPS). Simon et al. [118] investigated the production possibilities of floating houses through a highly automated construction and production planning framework. None of these structural studies, however, addresses the energy system design problem: coupled thermal-electrical modelling for structures with non-standard heat-transfer characteristics, stochastic consumption estimation from appliance-level specifications, and joint PV-storage optimisation under physical constraints such as weight limitations.

This chapter addresses these gaps. Section 4.2 presents the dichotomous search-based string optimisation (DSSO) algorithm, a consumption-aware method for multi-string PV energy systems under feed-in restrictions, with per-string logarithmic convergence guarantees for unimodal objectives, and evaluates it on a Hungarian Small-to-Medium Enterprise (SME) dataset across four surface configurations and two tariff structures. Section 4.3 develops the energy system design methodology for semi-autonomous structures, comprising the Stochastic Consumption Estimation Algorithm (SCEA), a coupled thermal-electrical system model, and a joint PV-storage optimisation approach evaluated in a year-long simulation of a semi-autonomous structure design (the floating-home case study). The formal thesis statements for Thesis III are gathered in Section 5.3, alongside those of Thesis I and Thesis II.

4.2 Consumption-aware dichotomous search-based string optimisation of multi-string photovoltaic systems

4.2.1 Motivation and context

As discussed in the chapter introduction (Section 4.1), industrial prosumers increasingly adopt PV energy systems to reduce energy costs and carbon emissions. However, many distribution system operators prohibit injecting excess energy into the public grid for industrial prosumers, owing to concerns about excessive PV capacity in the energy mix [7]. Under such feed-in restrictions, surplus PV generation is curtailed by feedback protection, representing energy lost to curtailment (E_{lost}). Traditional design methods fall short in this setting. The SYCM method relies only on annual energy aggregates and ignores the temporal mismatch between generation and consumption patterns, while the E_{lost} limit method compares simulated generation against the site's consumption profile but only caps curtailment, without optimising an economic objective. The two methods lead, respectively, to oversized systems with high curtailment and to undersized systems that inadequately reduce grid reliance. In the remainder of this chapter, design methods that exploit the temporal coupling between generation and consumption at high resolution (typically a 15-minute energy-balance simulation over a full year) are referred to as *consumption-aware*; in this sense the SYCM method is consumption-blind, and the E_{lost} limit method, although consumption-aware, lacks an optimisation framework.

Various optimisation methods have been presented for renewable energy systems with diverse objective functions. Ning et al. [74] proposed a Building Information Modelling (BIM)-based automatic design method that minimises the cost-to-power ratio. Liu et al. [71] analysed grid-connected residential PV energy systems in Australia using carbon dioxide emissions and economic metrics. Diaf et al. [73] presented a techno-economic optimisation of a hybrid wind-PV system on Corsica Island using LCOE. Fan and Xia [72] proposed a multi-objective retrofitting model using NPV and payback period. Ranaweera et al. [119] addressed building retrofitting from energy and environmental perspectives. Mehrabankhomartash et al. [75] analysed energy storage paired with PV energy systems using economic metrics combined with individual customer damage functions to size the battery. Starke et al. [76] proposed a multi-objective optimisation methodology, based on a genetic algorithm, for the design and sizing of hybrid concentrated solar power and PV (CSP+PV) plants with thermal energy storage, demonstrated on a case study for a location in Chile. Several studies have used commercial tools and pre-defined workflows for system design. Al-Karaghouli and Kazmerski [77] presented a HOMER-based optimisation and life-cycle cost analysis for the PV energy system of a rural health clinic in southern Iraq. Siregar et al. [78] used the commercial PVSyst tool to optimise and simulate a PV energy system on a site with a large placement area. Hong et al. [79] developed a method for estimating rooftop PV potential by analysing the available rooftop area with Hillshade analysis on a geographical information system (GIS) basis. On the algorithmic side, the dichotomous (bisection) search of [114] sits in a family of classical one-dimensional optimisers for unimodal functions on a bounded interval, alongside golden-section search and Brent's method, all of which converge logarithmically in the search-interval width. The choice among them for PV string sizing is revisited in Section 4.2.5. Distinct from the surveyed academic and tool-based approaches, the proposed algorithm exposes the objective function as a black box (obj(P) is queried at points selected by the bisection rule), optimises a multi-string system as a sequence of independent single-string problems ordered by sensitivity to the objective, and carries a per-string per-pass logarithmic convergence bound for any objective whose unimodality in string DC output power is empirically verified for the chosen tariff structure (verified here for NPV under both the fixed and the day-ahead market (DAM) tariffs of Section 4.2.3).

4.2.2 Method

The proposed consumption-aware evaluation framework couples simulated PV generation data, obtained from the PVGIS tool [110], with year-long measured consumption data at 15-minute resolution. For each candidate system configuration, a year-long simulation computes four energy-balance-aware metrics: the energy consumed from the public grid (E_{grid}), the energy curtailed by feedback protection (E_{lost}), the self-reliance ratio (μ_{sr}, the fraction of time during which the site requires no grid power), and the CO_2 emissions associated with grid consumption. The CO_2 dataset was constructed from hourly energy-mix data published by ENTSO-E [111], matched with production-type-specific emission factors from the IPCC report [112] and the UNECE life-cycle assessment of electricity generation options [113].

The DSSO algorithm optimises each string of a multi-string PV energy system separately, based on the dichotomous search principle. The recommended workflow for using the algorithm is shown in Figure 4.1. For a given string with feasible DC output power in the interval $[P^{\text{str}}_{\text{min}}, P^{\text{str}}_{\text{max}}]$, the algorithm maximises an arbitrary unimodal objective function $\text{obj}(P)$ by iteratively bisecting the interval. At each step, a differential is computed at the midpoint; since the objective is unimodal, the sign of the differential indicates which half of the interval contains the maximum, and the other half is discarded. The process repeats until the interval width falls below the convergence tolerance ε. Throughout this work, ε is set equal to the maximum DC output power of a single PV panel, the natural decision granularity for string sizing since PV panels are installed in integer counts; any string size returned by DSSO is therefore meaningful only up to that quantum, and a ε smaller than one panel would refine the search interval below the physically realisable resolution. The string sizes reported in this chapter are the unrounded bisection results; at installation they are rounded to the nearest whole panel. The unimodality assumption is treated as an empirical condition rather than a mathematical hypothesis: prior to applying DSSO to any new objective, the objective is swept on a fine 1D grid of string powers across $[P^{\text{str}}_{\text{min}}, P^{\text{str}}_{\text{max}}]$ to confirm that a single maximum exists. For NPV under both the fixed and the DAM tariffs of Section 4.2.3, this sweep revealed a single optimum and no local plateaus.

The order in which strings are optimised is determined by the sensitivity of the objective function to each string's output power, measured as the dichotomous differential of the objective evaluated at the midpoint of that string's search interval. For multi-string systems the algorithm operates as block coordinate descent: each string is bisected to its individual optimum holding the other strings fixed, then the loop revisits the strings in the same sensitivity order until no string size changes by more than ε between successive passes. This is a coordinate-descent procedure rather than a joint optimiser; for non-separable objectives such as NPV over a multi-string system (where curtailment couples the strings), the fixed point reached is a local optimum with respect to single-string perturbations, not necessarily the joint global optimum. The per-pass, per-string convergence rate is $O(\log_2((P^{\text{str}}_{\text{max}} - P^{\text{str}}_{\text{min}})/\varepsilon))$ iterations. The single-string DSSO algorithm is summarised in Algorithm 6.

4.2.3 Experimental methodology

The DSSO algorithm was evaluated using a real-world consumption dataset from a Hungarian SME. The year-long dataset had a 15-minute time resolution, with a sum yearly consumption of 322.62 MWh, an average input power of 2.3 kW, and a standard deviation of 1.77 kW.

Four PV surface configurations were tested: south-facing ground-mounted, south-facing roof-mounted, split east-west ground-mounted, and split east-west roof-mounted. Three design approaches were compared: the SYCM method, the E_{lost} limit method (targeting approximately 10% curtailment), and the proposed DSSO method. The DSSO objective was selected as the NPV of the investment, with a unit cost of 750 EUR/kWp (the average unit cost of installing a utility-scale PV system in Hungary in 2024) and a discount rate

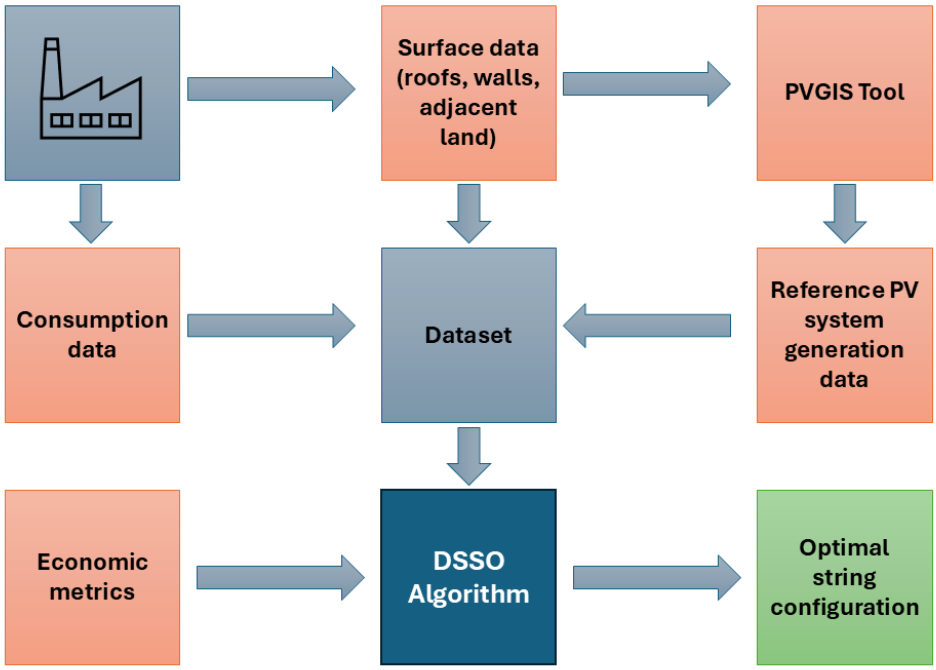


Figure 4.1: The recommended workflow for the DSSO algorithm. Surface data and consumption data are combined with reference PV generation data from the PVGIS tool to form the dataset, which is then optimised by the DSSO algorithm using economic metrics to produce the optimal string configuration.

$d = 5\%$. The NPV was evaluated from the annual energy-cost saving relative to the no-PV baseline, discounted at d over the project horizon, with module degradation, O&M costs, tariff escalation, and residual value excluded. Two tariff structures were tested: a fixed energy cost of 190 EUR/MWh and the DAM prices of the Hungarian Power Exchange (HUPX).

4.2.4 Results

The system sizes produced by each design method are reported in Table 4.1. The yearly aggregated energy-balance metrics for all configurations are summarised in Table 4.2. For the south ground-mounted surface with the fixed tariff, the DSSO-designed system required 107.8 kWp of PV panels and reached $\mu_{sr} = 22.35\%$, while the E_{lost} limit system required 20 kWp and reached $\mu_{sr} = 6.895\%$ and the SYCM system required 240 kWp. Across the four surfaces and the fixed tariff, the DSSO method produced systems sized at 37–45% of the corresponding SYCM systems. Compared with the E_{lost} limit method, the DSSO (fixed tariff) systems exhibited 14.8–17.5% lower grid consumption (E_{grid}), self-reliance ratios approximately three times higher, and 12.9–15% lower CO_2 emissions. Under the DAM tariff, the DSSO method produced smaller systems than under the fixed tariff. The east-west configurations produced asymmetric string sizes (see Table 4.1).

4.2.5 Discussion

The numerical results in Tables 4.1 and 4.2 confirm that the DSSO method addresses the gap identified in Section 4.2.1: the algorithm provides a principled optimisation framework for prosumers who cannot inject energy into the grid. Compared with the E_{lost} limit method, the fixed-tariff DSSO systems achieve substantially lower grid consumption and approximately three times higher self-reliance, while also reducing CO_2 emissions, despite operating on a

Algorithm 6: The DSSO algorithm for single-string PV sizing. Dichotomous (bisection) search maximises the unimodal objective function obj(P) by discarding half of the power interval at each step, converging in $O(\log_2((P^{\text{str}}_{\text{max}} - P^{\text{str}}_{\text{min}})/\varepsilon))$ iterations to the optimal installed peak power $P^{\text{str}}_{\text{opt}}$ of the string.

Input: $P^{\text{str}}_{\text{max}}$ max installed DC peak power (kWp)
$P^{\text{str}}_{\text{min}}$ min installed DC peak power (kWp)
ε convergence tolerance
obj(P) unimodal objective function
Output: $P^{\text{str}}_{\text{opt}}$ optimal installed DC peak power (kWp) of the string

// Step 1: Test boundaries
$\Delta P(P^{\text{str}}_{\text{min}}) \leftarrow \text{obj}(P^{\text{str}}_{\text{min}}) - \text{obj}(P^{\text{str}}_{\text{min}} + \varepsilon)$;
if $\Delta P(P^{\text{str}}_{\text{min}}) \geq 0$ **then return** $P^{\text{str}}_{\text{min}}$;
$\Delta P(P^{\text{str}}_{\text{max}}) \leftarrow \text{obj}(P^{\text{str}}_{\text{max}}) - \text{obj}(P^{\text{str}}_{\text{max}} - \varepsilon)$;
if $\Delta P(P^{\text{str}}_{\text{max}}) \geq 0$ **then return** $P^{\text{str}}_{\text{max}}$;
// Step 2: Bisection
$P_{\text{low}} \leftarrow P^{\text{str}}_{\text{min}}$; $P_{\text{high}} \leftarrow P^{\text{str}}_{\text{max}}$;
$P_{\text{mid}} \leftarrow (P^{\text{str}}_{\text{max}} - P^{\text{str}}_{\text{min}})/2 + P^{\text{str}}_{\text{min}}$;
while $P_{\text{high}} - P_{\text{low}} \geq \varepsilon$ **do**
  $\Delta P(P_{\text{mid}}) \leftarrow \text{obj}(P_{\text{mid}}) - \text{obj}(P_{\text{mid}} - \varepsilon)$;
  if $\Delta P(P_{\text{mid}}) < 0$ **then** $P_{\text{high}} \leftarrow P_{\text{mid}}$;
  else $P_{\text{low}} \leftarrow P_{\text{mid}}$;
  $P_{\text{mid}} \leftarrow (P_{\text{high}} - P_{\text{low}})/2 + P_{\text{low}}$;
return P_{low}

Table 4.1: System sizes (in kWp) designed with the SYCM method, the E_{lost} limit method, and the proposed DSSO method for each surface configuration. East-west configurations show string 1 + string 2 sizes. Under the fixed tariff the DSSO systems are sized at 37–45% of the corresponding SYCM systems, and the DAM tariff yields smaller systems still.

Design method	S. ground	S. roof	E-W ground	E-W roof
SYCM	240	259	155+155	148+148
E_{lost} limit	20	21.5	15+15	12+12
DSSO (fixed)	107.8	108.2	49.6+77.7	25.8+82.3
DSSO (exchange)	67	57.2	21.4+33.4	8.8+36.2

smaller footprint than the SYCM systems. The asymmetric string sizing observed in the east-west configurations reflects the value of small departures from symmetry under the chosen NPV objective: the algorithm prioritises the more efficient surface and adjusts the less efficient one to balance overproduction against curtailment. The smaller systems produced under the DAM tariff reflect the additional value that can be captured through time-of-use pricing, which the SYCM and E_{lost} limit methods cannot exploit.

The key strengths of the DSSO algorithm are its applicability to any objective function that is unimodal in string DC output power (economic or sustainability-based), its support for any tariff structure that preserves this unimodality (including the fixed and DAM tariffs tested here), its logarithmic per-string per-pass convergence, and its ability to produce consumption-aware designs that balance system size, grid reliance, and environmental impact. Tariffs with discontinuities, such as demand charges tied to peak power or block-

Table 4.2: Yearly aggregated energy metrics for systems designed with the SYCM, E_{lost} limit, and DSSO methods across the four surface configurations. E_{grid}: grid consumption (MWh); E_{lost}: curtailed energy (MWh); μ_{sr}: self-reliance ratio (%); CO_2: yearly emissions (t). Relative to the E_{lost} limit method, the fixed-tariff DSSO systems roughly triple the self-reliance ratio and lower both grid consumption and CO_2 emissions.

Metric	Method	S. ground	S. roof	E-W ground	E-W roof
E_{grid} [MWh]	SYCM	219.00	217.35	205.29	215.16
	E_{lost} limit	298.72	298.55	294.77	299.01
	DSSO (fixed)	246.49	249.38	244.85	254.71
	DSSO (exchange)	265.08	273.21	278.34	284.38
E_{lost} [MWh]	SYCM	215.9	217.4	203.3	215.2
	E_{lost} limit	2.675	2.701	3.182	2.502
	DSSO (fixed)	67.08	61.47	54.30	50.14
	DSSO (exchange)	31.47	21.80	12.57	10.94
μ_{sr} [%]	SYCM	30.36	30.79	32.27	30.79
	E_{lost} limit	6.895	7.174	8.253	7.026
	DSSO (fixed)	22.35	22.00	22.00	20.62
	DSSO (exchange)	17.41	15.80	13.90	12.59
CO_2 [t]	SYCM	31.60	31.24	29.44	30.83
	E_{lost} limit	40.88	40.84	40.33	40.87
	DSSO (fixed)	34.87	35.14	34.29	35.59
	DSSO (exchange)	37.03	37.93	38.33	39.11

rate structures with stepwise discounts, may break the unimodality of NPV in string power and would therefore require either an objective reformulation or a search method that does not rely on unimodality before DSSO is applied. The choice of dichotomous bisection rather than golden-section search or Brent's method was made deliberately. Although golden-section achieves the same $O(\log)$ convergence rate with one function evaluation per iteration instead of two, the dichotomous differential evaluated at the midpoint admits a direct physical interpretation: the change in the objective when string power is altered by one panel's worth of DC capacity, the natural decision granularity for PV string sizing. This interpretability outweighs the constant-factor evaluation saving offered by golden-section or Brent's method in the regime where the optimisation budget is dominated by the per-iteration year-long simulation rather than by the number of iterations. A limitation is the unimodality assumption on the objective function, which excludes multi-modal objectives such as those arising in joint sizing problems with BESS under multiple competing targets. A second limitation of the evaluation is that DSSO is benchmarked only against two industry rules of thumb (SYCM and the E_{lost} limit method) and not against published optimisation-based sizing methodologies such as HOMER, the PVSyst optimiser, mixed-integer linear programming (MILP)-based sizing, or genetic-algorithm-based hybrid-system design; an objective-by-objective comparison against those methods is reserved for future work. A further assumption of the DSSO framework is that measured consumption data at high temporal resolution is available; for cases where such data does not exist (for example, when designing the energy system of a new structure), a synthetic load profile must first be generated. This requirement motivates the SCEA presented in the following section.

4.3 Design methodology for semi-autonomous electric energy systems

4.3.1 Motivation and context

As noted in the chapter introduction (Section 4.1), semi-autonomous structures that rely primarily on local PV generation and BESS require a joint design methodology that accounts for the coupling between generation, storage, thermal loads, and physical constraints. Unlike grid-connected prosumers, these systems must achieve high self-reliance with limited or absent grid backup, making isolated component-level optimisation insufficient.

The floating-home case study considered in this section exemplifies this class of semi-autonomous structure and presents several additional constraints that make it a challenging validation case: (i) the structure can be rotated within constraints, enabling PV orientation optimisation; (ii) the insulation and sun exposure differ from conventional buildings due to the water-surrounded environment, affecting the thermal balance; (iii) the physical weight capacity of the platform limits BESS sizing; and (iv) the limited or absent grid connection demands the highest possible self-reliance ratio.

Existing research on the energy system of semi-autonomous structures provides only partial coverage of this design problem. Smart-home and microgrid reviews [80] establish the broader taxonomy of energy management strategies. Akinyele et al. [81] demonstrated the feasibility of an autonomous home through a proof-of-concept case study. Aguilar et al. [115] addressed the consumption-control side, presenting an autonomous cycle of data-analysis tasks for scheduling controllable load appliances so that non-controllable renewable generation is utilised more fully. Tipantuna and Hesselbach [116] addressed the system-level side with an IoT-enabled proposal for adaptive self-powered renewable-energy management in home systems. For water-surrounded structures, the literature is dominated by structural and architectural studies. Ishaque et al. [84] described a structural design for a low-cost floating house in which safety was the primary criterion. Cahyono [117] reported a stability test of a floating-house platform with a lightweight concrete cover made of EPS, the construction style closest to the case study considered here. Simon et al. [118] investigated the production possibilities of floating houses through a highly automated construction and production planning framework. None of these contributions, however, addresses the joint design problem of the energy system: stochastic consumption estimation from appliance-level specifications, coupled thermal-electrical modelling for structures with non-standard heat-transfer characteristics, and joint PV-storage optimisation under physical constraints such as weight limitations.

4.3.2 Method

The proposed methodology addresses the joint design of PV generation and BESS for semi-autonomous structures. The workflow is illustrated in Figure 4.2. It consists of three main stages. First, the combined consumption of all appliances and the output power of the PV system are estimated for a full year. Second, this dataset is extended with electricity prices from the DAM of an energy exchange and used as input to a year-long simulation of the system model. Third, the simulation results are aggregated and visualised. The key difference from traditional workflows is the optimisation order: traditional methods optimise the PV system for maximum output first, then estimate consumption, whereas the proposed methodology optimises the PV system with the entire energy system in mind, since battery-enhanced systems can sometimes utilise more energy with longer charging times rather than higher input power.

The PV output of each design is estimated using the API of the industry-standard PVGIS tool [110], which returns hourly PV output and ambient-temperature time series for any given surface orientation. For structures whose orientation is not architecturally fixed (as is the case for the floating-home case study considered later, which can be rotated within

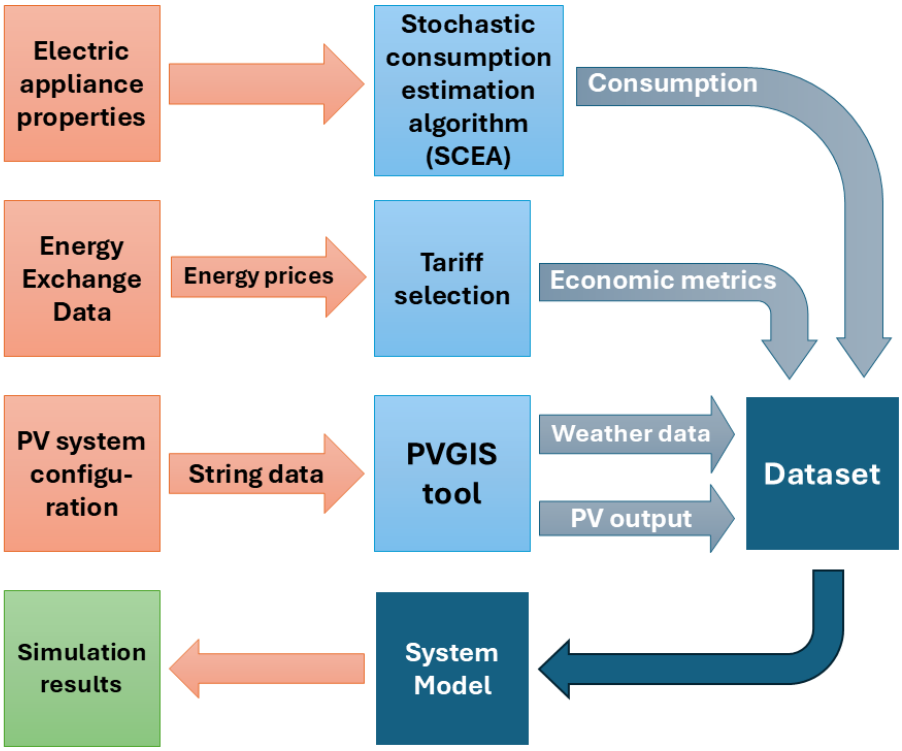


Figure 4.2: Block diagram of the simulation workflow for the semi-autonomous energy system design methodology. The SCEA generates load profiles from appliance properties, which are combined with PV output data, weather data, and electricity prices to form the input dataset for the system model.

mooring constraints), the methodology supports a two-variable orientation optimisation: the tilt θ of the roof-mounted string and the azimuth ϕ of the structure as a whole. The orientation (θ^*, ϕ^*) is obtained by alternating single-variable line searches in θ and ϕ until a local maximum of the annual PV yield is reached, followed by a dual-variable sensitivity sweep in the (θ, ϕ) neighbourhood to confirm local optimality within the sampled neighbourhood. For structures whose orientation is fixed by the architecture or by site constraints, (θ, ϕ) is taken as given and PVGIS is queried once per surface; this is the regime used in the case study, where the roof tilt and the east/west wall orientations are determined by the existing structural design.

The SCEA synthesises year-long load profiles from appliance-level properties and probabilistic usage patterns. For each appliance, the inputs are: the maximum power $P^{\text{app}}_{\text{max}}$, the minimum power $P^{\text{app}}_{\text{min}}$, the fill ratio Λ (the fraction of time during which the appliance consumes high power while on), and a 24-element usage weight profile $\mathbf{W}_{\text{profile}}$ of relative hourly usage weights (normalised to sum to one). The algorithm proceeds in three steps, summarised in Algorithm 7: (i) initialise an array of random states for each time slot; (ii) normalise the weight profile to produce probabilities; and (iii) compute the energy consumed in each time slot using Equation (4.1):

$$E_{\text{used}}[idx] = \frac{\Delta t}{60} \cdot \left[\mathbf{W}_{\text{profile}}[\lfloor idx \cdot \Delta t/60 \rfloor] \cdot u[idx] \cdot \Lambda \cdot \left(P^{\text{app}}_{\text{max}} - P^{\text{app}}_{\text{min}}\right) + P^{\text{app}}_{\text{min}}\right], \tag{4.1}$$

where $u[idx] \sim \mathcal{U}(0,1)$ is a uniform random variable for each time slot, and the hour-of-day index $\lfloor idx \cdot \Delta t/60 \rfloor \in [0, 23]$ maps each slot of width Δt minutes to its containing hour. The leading factor $\Delta t/60$ converts the slot-averaged power within the brackets (in kW) to a per-slot energy (in kWh). The output E_{used} is a year-long time series at the specified resolution Δt (15 minutes in the case study, giving $24 \cdot 60/\Delta t = 96$ slots per day). Individual appliance profiles are summed to obtain the total site load.

Algorithm 7: The SCEA for synthetic load-profile generation. For each time slot a random draw is scaled by the appliance's hourly usage weight, fill ratio, and power range; the product of the draw and the fill ratio acts as a continuous usage-intensity factor rather than an on/off state. The procedure returns one day's profile for one appliance; it is invoked independently for each of the 365 days and each appliance, and the results are summed slot-wise to give the year-long load profile at resolution Δt.

Input: Δt output resolution (minutes)
$P^{\text{app}}_{\text{max}}$ max appliance power
$P^{\text{app}}_{\text{min}}$ min appliance power
Λ fill ratio
$\mathbf{W}_{\text{profile}}$ 24-element hourly weight profile
Output: E_{used} simulated consumption array for one day

// Step 1: Initialise random states
for $idx = 0$ **to** $24 \cdot 60/\Delta t - 1$ **do**
  $u[idx] \leftarrow \text{random}([0, 1])$;
// Step 2: Normalise the weight profile
$S \leftarrow \text{SUM}(\mathbf{W}_{\text{profile}})$;
for $idx = 0$ **to** 23 **do**
  $\mathbf{W}_{\text{profile}}[idx] \leftarrow \mathbf{W}_{\text{profile}}[idx] \,/\, S$;
// Step 3: Compute the load profile
for $idx = 0$ **to** $24 \cdot 60/\Delta t - 1$ **do**
  $E_{\text{used}}[idx] \leftarrow \frac{\Delta t}{60} \cdot \left[\mathbf{W}_{\text{profile}}[\lfloor idx \cdot \Delta t/60 \rfloor] \cdot u[idx] \cdot \Lambda \cdot (P^{\text{app}}_{\text{max}} - P^{\text{app}}_{\text{min}}) + P^{\text{app}}_{\text{min}}\right]$;
return E_{used}

The SCEA is calibrated for aggregate annual consumption rather than for the fine structure of individual appliance traces, and four scope-defining choices follow directly from this stance. First, the per-slot random draw $u[idx]$ is independent across slots, so the algorithm does not reproduce the within-appliance burst structure of real usage (a microwave on at time t and at $t + \Delta t$ are independent events in the model, whereas in reality they are highly correlated); quantities sensitive to burst structure such as peak demand or distribution-network coincidence factors must therefore not be read off the SCEA output. Second, the random draws are independent across appliances, so positively correlated activities (meal preparation, evening occupancy) are not captured and the variance of the aggregate load is the sum of per-appliance variances rather than the larger correlated value. Third, the 24-element weight profile $\mathbf{W}_{\text{profile}}$ is held identical across all days of the year, with no weekday/weekend distinction, no monthly seasonality, and no holiday effect. Fourth, the algorithm has no explicit occupancy model: the home is assumed to be perpetually accessible to all appliances. These four reductions are acceptable because the SCEA output is consumed only by a year-long energy-balance simulation whose figures of merit (μ_{sr}, E_{lost}, grid-consumption decrease) are slot-aggregated and depend almost entirely on the annual energy total and the mean intra-day load shape, both of which the SCEA reproduces by construction (the total exactly, the shape in expectation). Within the workflow, the SCEA is therefore positioned as a deliberately simple stratified-sampling baseline calibrated to annual energy aggregates rather than as an attempt to compete with the richer published synthetic-load-generation literature; load generators with higher per-appliance realism, such as the CREST stochastic high-resolution household-electricity model or BREDEM-style domestic-energy models, can be substituted for SCEA as plug-in replacements at the

workflow's load-input stage whenever higher load-shape fidelity is required.

The coupled thermal-electrical model captures the interactions between the air-conditioning system, the BESS, the PV generation, and the structure's thermal dynamics. The thermal subsystem models the air-conditioning system using a temperature-dependent coefficient of performance (COP). The thermal power injected into the structure by the air-conditioning system is given by Equation (4.2):

$$Q_{\mathrm{ac}} = s_{\mathrm{ac}} \cdot \left(\mathrm{COP}_{\mathrm{nom}} + k_{\mathrm{accop}} \cdot (T_{\mathrm{out}} - T_{\mathrm{target}})\right) \cdot P_{\mathrm{ac}}, \tag{4.2}$$

where $\mathrm{COP}_{\mathrm{nom}}$ is the nominal COP, k_{accop} is the temperature-difference correction coefficient, T_{out} is the outside temperature, T_{target} is the inside target temperature, $P_{\mathrm{ac}} \geq 0$ is the air-conditioning electrical input power, and $s_{\mathrm{ac}} \in \{+1, -1\}$ is the operating-mode indicator. By convention, Q_{ac} is the signed net thermal flow into the interior: $s_{\mathrm{ac}} = +1$ in heating mode gives $Q_{\mathrm{ac}} > 0$, and $s_{\mathrm{ac}} = -1$ in cooling mode gives $Q_{\mathrm{ac}} < 0$. Keeping P_{ac} non-negative lets the same quantity enter the electrical balance of Equation (4.5) as a load in both modes.

Heat dissipation between the structure and the outside environment is modelled by Equation (4.3):

$$Q_{\mathrm{diss}} = k_{\mathrm{diss}} \cdot (T_{\mathrm{out}} - T_{\mathrm{inside}}), \tag{4.3}$$

where k_{diss} is the dissipation coefficient and T_{inside} is the inside temperature. Additional heat sources include solar radiation on the structure's surfaces (modelled with a coefficient k_{PV}) and waste heat from battery power electronics. After simplification, the total heat introduced into the interior is given by Equation (4.4):

$$Q_{\mathrm{inside}} = Q_{\mathrm{ac}} + (k_{\mathrm{peb}} + k_{\mathrm{chem}}) \cdot (P_{\mathrm{bin}} + P_{\mathrm{bout}}) + k_{\mathrm{PV}} \cdot P_{\mathrm{PV}} - Q_{\mathrm{vent}} + Q_{\mathrm{app}} + Q_{\mathrm{diss}}, \tag{4.4}$$

where k_{peb} and k_{chem} are the power-electronics and battery-chemistry loss coefficients, P_{bin} and P_{bout} are the non-negative charging and discharging powers (so the battery loss term is a heat source in both directions), P_{PV} is the PV output power, Q_{vent} is the ventilation heat loss, and Q_{app} is the appliance heating power. The inside temperature change over one slot is $\Delta T_{\mathrm{inside}} = Q_{\mathrm{inside}} \cdot (\Delta t/60)/C_{\mathrm{inside}}$, where C_{inside} is the heat capacity of the interior and $\Delta t/60$ is the slot length in hours.

The thermal model deliberately collapses the structure to a single interior temperature node, and five reductions concentrate in equations (4.2) to (4.4). First, the interior is represented as a single lumped thermal mass C_{inside}: wall capacitance, furniture, and the water-side mass of the floating platform are aggregated rather than modelled as a 2R1C or 3R2C network, so thermal lag from the wall is aliased into the dissipation coefficient k_{diss}. Second, k_{diss} in Equation (4.3) aggregates the air-side and water-side heat transfers into a single value referenced to T_{out}, on the assumption that the platform's underside insulation dominates the water-side transfer; this assumption is reasonable for a well-insulated foam-based floating module but should be re-evaluated for thinly insulated platforms. Third, the solar gain is modelled as $k_{\mathrm{PV}} \cdot P_{\mathrm{PV}}$, an effective lumped coefficient calibrated against the PVGIS-derived envelope irradiance rather than a physical fraction of the PV output: k_{PV} should not be interpreted as a fixed percentage of P_{PV} and is not transferable to structures with different envelope-to-PV area ratios. Fourth, the COP is linearised in $(T_{\mathrm{out}} - T_{\mathrm{target}})$ via Equation (4.2), with the linearisation assumed valid over the −5 to +35 °C outdoor range taken for the Hungarian case-study site; refrigerant-cycle saturation and defrost cycles fall outside this range and would require a piecewise model. Fifth, the air-conditioning input power P_{ac} is treated as exogenous in Equation (4.2) but is in fact set by a deadband thermostat in the simulation; the control law $P_{\mathrm{ac}} = \mathrm{thermostat}(T_{\mathrm{inside}}; T_{\mathrm{target}} \pm 1\,\mathrm{K}, \overline{P_{\mathrm{ac}}})$ closes the thermal loop and bounds the air-conditioning draw. Domestic hot water, latent heat, humidity, and interior thermal stratification are out of scope and are deferred to a richer model.

The electrical subsystem satisfies the slot-averaged power balance of Equation (4.5):

$$P_{\mathrm{grid}} + P_{\mathrm{PV}} + P_{\mathrm{bout}} = P_{\mathrm{app}} + P_{\mathrm{ac}} + P_{\mathrm{bin}} + P_{\mathrm{PV,lost}}, \tag{4.5}$$

where P_{grid} is the power consumed from the grid, P_{app} is the appliance load, P_{bin} and P_{bout} are the battery charge and discharge powers, and $P_{\text{PV,lost}}$ is the unutilised PV power. All quantities denoted $P.$ are slot-averaged powers in kW; the corresponding per-slot energies in kWh are obtained by multiplying by the time-step length $\Delta t/60$ in hours. The battery dynamics are governed by Equation (4.6):

$$\Delta E_{\text{st}} = \left[P_{\text{bin}} \cdot (1 - k_{\text{peb}} - k_{\text{chem}}) - P_{\text{bout}} \cdot (1 + k_{\text{peb}} + k_{\text{chem}})\right] \cdot \frac{\Delta t}{60}, \tag{4.6}$$

subject to the constraints $P_{\text{bin}} = 0$ and $\Delta E_{\text{st}} = 0$ for charging when $E_{\text{st}} \geq C_{\text{st}}$ (fully charged), and $P_{\text{bout}} = 0$ and $\Delta E_{\text{st}} = 0$ for discharging when $E_{\text{st}} \leq 0$ (fully discharged); the charging and discharging powers are bounded by the converter ratings, $0 \leq P_{\text{bin}} \leq P_{\text{bin,max}}$ and $0 \leq P_{\text{bout}} \leq P_{\text{bout,max}}$, and are mutually exclusive within a slot, $P_{\text{bin}} \cdot P_{\text{bout}} = 0$. Within each slot the dispatch follows a fixed priority: PV generation first serves the load; any surplus charges the battery up to $P_{\text{bin,max}}$ and the remaining free capacity, and whatever is left is counted as unutilised ($P_{\text{PV,lost}}$); a deficit is met from the battery up to $P_{\text{bout,max}}$ and then from the grid within the connection limit. The BESS capacity C_{st} is constrained by the physical limitations of the structure (e.g., weight capacity in the case of water-surrounded platforms).

The electrical subsystem reduces the storage and conversion chain to a single linearised reservoir, and four scope-defining choices concentrate in Equations (4.5) and (4.6). First, the loss coefficients k_{peb} and k_{chem} play a dual role: they appear in the battery dynamics as electrical-loss fractions and in Equation (4.4) as thermal-injection fractions, on the assumption that the entirety of the electrical losses is dissipated as heat into the structure; this identification holds for indoor-mounted electronics but breaks for outdoor-vented inverters, in which case the thermal contribution should be removed from Equation (4.4). Second, the loss coefficients are held constant with respect to state of charge (SoC), C-rate, battery temperature, and calendar age, so the model does not capture the SoC-dependent efficiency tail of lithium chemistries or the temperature-dependent loss surge in cold operation; for a case study at near-room ambient, this is a reasonable first approximation. Third, the storage window is the full nominal range $E_{\text{st}} \in [0, C_{\text{st}}]$ with no working-window margins; real lithium chemistries are usually limited to roughly 10–90% SoC for cycle-life reasons, which means that the reported unutilised energy E_{lost} is a lower-bound estimate relative to a working-window-aware dispatch, which would absorb less surplus, and the reported self-reliance is correspondingly optimistic. Fourth, the PV output P_{PV} enters Equation (4.5) as a known exogenous quantity sourced from the PVGIS tool with its default 14% combined system-loss factor; separate inverter modelling, partial-load efficiency curves, soiling, and module-temperature derating beyond what PVGIS already includes are out of scope.

The methodology frames the joint PV–BESS design problem as a constrained discrete optimisation. Let $\mathcal{D}$ denote the finite catalogue of feasible (PV-configuration, BESS-configuration) pairs c available to the designer. The design problem is to find

$$c^* = \arg\min_{c \in \mathcal{D}_\tau} E_{\text{lost}}(c), \qquad \mathcal{D}_\tau = \left\{ c \in \mathcal{D} \;:\; \mu_{\text{sr}}(c) \geq \max_{c' \in \mathcal{D}} \mu_{\text{sr}}(c') - \tau \right\}, \tag{4.7}$$

that is, among the configurations whose self-reliance is within a tolerance τ of the best attainable, the one that leaves the least generation unutilised is selected; $\tau = 2$ percentage points throughout this work. The optimisation is subject to three physical constraints: (i) zero feed-in, so any excess PV generation that cannot be consumed or stored within the same time step remains unutilised and is counted as E_{lost}; (ii) a fixed grid-connection limit $P_{\text{grid}} \leq P_{\text{conn,max}}$ that bounds the instantaneous grid draw at the point of common coupling; and (iii) a platform-imposed upper bound $C_{\text{st}} \leq C_{\text{st,max}}$ on the installed battery capacity, set by the structure's weight allowance. Because $\mathcal{D}$ is finite and small, the optimum is identified

by exhaustive enumeration: every candidate $c \in \mathcal{D}$ is evaluated by a full-year simulation under Equations (4.2) to (4.6), infeasible candidates are discarded, the near-optimal set $\mathcal{D}_\tau$ is formed from the maximum attainable μ_{sr}, and c^* is the member of $\mathcal{D}_\tau$ with the lowest E_{lost}. The lexicographic form encodes an engineering and economic preference: once self-reliance is within τ of the best attainable, further generation that cannot be consumed or stored buys nothing and only adds installed capacity.

A key simplification in the case study was the decoupling of the thermal and electrical models: since the battery waste heat was vented outside rather than recovered for heating, the battery loss term $(k_{\text{peb}}+k_{\text{chem}})\cdot(P_{\text{bin}}+P_{\text{bout}})$ in Equation (4.4) is zero for this structure, and the two models could be processed sequentially, reducing computational complexity while maintaining accuracy. The year-long simulation starts with an empty battery ($E_{\text{st}}(0) = 0$) and imposes no terminal state-of-charge condition.

4.3.3 Experimental methodology

The methodology was evaluated in a year-long simulation of the floating-home case study, the design of a real-world semi-autonomous structure combining PV generation and BESS under tight physical constraints. The structure was designed for up to four persons, with a 16 A single-phase grid connection giving $P_{\text{conn,max}} = 3.7$ kW (no feed-in permitted) and a composite concrete-coated foam-based platform. Three PV surfaces were available: a roof supporting up to 17.6 kWp at 40° tilt, and two vertical walls (east: 4.63 kWp; west: 4.42 kWp) at 90° tilt, for a full installable capacity of 26.65 kWp. The sum yearly appliance consumption (excluding heating and cooling) was 3.872 MWh; the sum yearly heating and cooling demand was 3 MWh of delivered thermal energy.

12 scenarios were evaluated, formed as the Cartesian product of four PV configurations and three storage configurations. The four PV configurations were: full PV system (all surfaces), half PV system (half of each surface), roof only, and east-west walls only. The three storage configurations were: no storage; 10 kWh capacity with 5 kW input/output power; and 20 kWh capacity with 10 kW input/output power. Both batteries had a round-trip efficiency of 90%. The simulation ran with 15-minute resolution over a full calendar year. Three metrics were aggregated for each scenario: the percentage decrease in grid consumption, the self-reliance ratio μ_{sr}, and the unutilised (curtailed) PV energy.

4.3.4 Results

The percentage decrease in grid consumption for each PV-storage combination is reported in Table 4.3. The self-reliance ratio is reported in Table 4.4, and the unutilised PV energy is reported in Table 4.5. The full PV system with the largest storage (C_{st} = 20 kWh) achieved the highest headline metrics: an 86.2% reduction in grid consumption (Table 4.3) and 91.43% self-reliance (Table 4.4). For the same configuration, however, 22.14 MWh/year of PV energy remained unutilised (Table 4.5), the largest such total among the four PV configurations at this battery size. The half PV system with the largest storage reached 78.5% grid reduction with 8.28 MWh/year unutilised; the corresponding figures were 83.7% and 16.53 MWh/year for the roof-only system, and 56.4% and 1.35 MWh/year for the east-west-only system.

Table 4.3: Percentage decrease in grid consumption for each PV-storage combination. All values are percentages. The full PV system with the largest storage (20 kWh) gives the greatest reduction (86.2%), and the reduction grows with storage size for every configuration.

Storage configuration	Full PV	Half PV	Roof only	E-W only
No storage	27.9	25.6	26.3	23.0
10 kWh / 5 kW	67.1	62.2	60.8	50.3
20 kWh / 10 kW	86.2	78.5	83.7	56.4

Table 4.4: Self-reliance ratio μ_{sr} (%) for each PV-storage combination. Self-reliance peaks at 91.43% for the full PV system with 20 kWh storage, with the roof-only system reaching almost the same value (89.59%).

Storage configuration	Full PV	Half PV	Roof only	E-W only
No storage	42.67	39.63	40.93	34.31
10 kWh / 5 kW	80.31	75.94	78.45	63.77
20 kWh / 10 kW	91.43	85.47	89.59	67.47

Table 4.5: Unutilised (curtailed) PV energy in MWh for each PV-storage combination. The full PV system leaves the most energy unutilised (22.14 MWh with the largest battery), and the amount falls as installed PV capacity is reduced, exposing the oversizing penalty under weight-limited storage.

Storage configuration	Full PV	Half PV	Roof only	E-W only
No storage	26.81	12.51	21.17	4.02
10 kWh / 5 kW	23.6	9.5	17.98	1.84
20 kWh / 10 kW	22.14	8.28	16.53	1.35

4.3.5 Discussion

The numerical results in Tables 4.3 to 4.5 confirm that the proposed methodology produces actionable design decisions for semi-autonomous energy systems and addresses the joint-design gap identified in Section 4.3.1. Applied to the case-study feasibility set, self-reliance is maximised by the full PV system with the 20 kWh / 10 kW BESS (91.43% in Table 4.4), so the near-optimal set $\mathcal{D}_\tau$ of Equation (4.7) comprises the configurations

reaching at least 89.43%. The roof-only system with the same battery falls inside that band at 89.59% self-reliance and an 83.7% grid-consumption reduction (against 86.2% for the full system), while leaving substantially less generation unutilised: 16.53 MWh/year against 22.14 MWh/year (Table 4.5). By Equation (4.7) the roof-only configuration is therefore the optimum c^*, and the full PV system is rejected as oversized for this structure. The SCEA algorithm (Algorithm 7) generated load profiles calibrated to the 3.872 MWh/year appliance-level total specified for the structure, providing a viable input when measured consumption data is unavailable. The coupled thermal-electrical model (Equations (4.2) to (4.6)), despite the simplifications adopted in the case study (decoupling of thermal and electrical subsystems through outside venting of battery waste heat), captured the interactions that determine system performance. The dominance of the roof surface as the most beneficial mounting location, despite the intuitive appeal of maximising total PV area by also using the vertical walls, follows directly from the storage constraint: under a weight-limited BESS capacity, additional generation that cannot be stored remains unutilised. This explains why the full PV system, which has the highest installed capacity, also has the largest figure in Table 4.5 (22.14 MWh/year with the largest battery). The same mechanism is expected to arise in other semi-autonomous structures with constrained storage capacity whenever PV generation is optimised in isolation, though confirming this would require applying the methodology to further instances. The half-system configuration offers a further compromise, achieving 78.5% grid-consumption reduction with the largest storage while leaving less energy unutilised (8.28 MWh/year). Although the full PV system with the largest storage attains the highest headline metrics, its additional generation cannot be put to use; accounting for this oversizing penalty, the roof-only configuration remains the more effective choice, consistent with the dominance of the roof as the most beneficial mounting surface.

A first limitation of the evaluation is that the methodology has been demonstrated on a single semi-autonomous structure; broader claims of generalisability would require its application to additional instances: other floating homes with different roof areas and platform weight allowances, land-based semi-autonomous structures such as off-grid cabins or stand-alone agricultural buildings, and structures situated in climate zones outside the temperate-Central-European range of the case study. A second limitation is that the validation relies on year-long simulation rather than on measurements from the completed structure; the headline self-reliance and grid-consumption-reduction figures are therefore conditional on the calibration of the thermal-electrical model, with the largest uncertainties carried by the dissipation coefficient k_{diss}, the solar-gain coefficient k_{PV}, the COP correction coefficient k_{accop}, and the combined battery loss term $k_{\text{peb}} + k_{\text{chem}}$. A formal parameter-sensitivity analysis was not part of the case study and is reserved for the post-construction validation. A third limitation is that the methodology was not benchmarked, on the same case study, against the traditional sequential design workflow that motivated it against which the joint procedure of Equation (4.7) is positioned as an improvement. Quantifying the joint-vs-sequential delta on the floating-home case study would require a separate simulation pass and is reserved for future work. Each of the three stages of the workflow (the SCEA, the coupled thermal-electrical model, and the discrete design procedure of Equation (4.7)) is in principle independently usable: SCEA can feed any external design procedure, the thermal-electrical model can validate any sizing decision, and the design procedure could in principle consume measured rather than synthesised consumption. The contribution claimed under Thesis III.2 is the integrated workflow, not the three stages in isolation; benchmarking each stage against off-the-shelf alternatives is a complementary direction for future work. Future work targets the validation of the simulation results on the finished semi-autonomous structure, the further simplification of the thermal-electrical model where accuracy permits, and the integration of an automatic system-level optimisation routine that combines the joint design methodology with the consumption-aware string optimisation introduced in Section 4.2.

Chapter 5
New Scientific Results

5.1 Thesis I: Day-ahead electricity price forecasting using machine learning

The contributions of Thesis I are summarised as follows.

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.

Related publications:

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]**

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.

Related publications:

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]**

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]**

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.

Related publications:

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]**

5.2 Thesis II: Energy trading strategy optimisation for battery-enhanced power plants

The contributions of Thesis II are summarised as follows.

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.

Related publications:

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]**

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.

Related publications:

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]**

5.3 Thesis III: Physical design optimisation of photovoltaic energy systems

The contributions of Thesis III are summarised as follows.

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 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_{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.

Related publications:

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 *2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET)*, pages 1–6, 2025.

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.

Related publications:

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 *2025 IEEE 29th International Conference on Intelligent Engineering Systems (INES)*, pages 143–148, 2025.

Chapter 6
Conclusion

I addressed three interconnected challenges along a decision-support chain for battery-enhanced renewable energy systems: forecasting day-ahead prices, optimising trading strategies with predicted prices, and designing the physical battery-enhanced photovoltaic (PV) energy systems that participate in the resulting markets. Together they form the prediction–control–design pipeline summarised in the graphical abstract (Figure 1.1). The three stages were independently validated against published data and case studies. Their integration into a single end-to-end workflow was an overarching goal of the research, which was partitioned into three interconnected research objectives.

The first objective (Objective 1) was the development of machine learning methods for accurate electricity price forecasting (EPF) on the day-ahead market (DAM), the upstream input on which any storage-aware trading strategy depends. Three contributions addressed this objective: a general-purpose stopping criterion, a hyperparameter-optimised Long Short-Term Memory (LSTM) model and training pipeline, and a systematic study of pre-trained Transformer foundation models with context length optimisation for direct multi-step forecasting.

In *Thesis I.1* I introduced the Correlation-Driven Stopping Criterion (CDSC), a neural network (NN) regularisation method that monitors the rolling Pearson correlation r between training and validation losses within a sliding window of size ω and terminates training once r has been below a threshold μ_{cdsc} for λ_{cdsc} epochs in total, not necessarily consecutive ones. Where early stopping reacts only to the validation minimum, the CDSC captures the joint dynamics of both curves and detects the onset of overfitting more precisely. Across four benchmarks (binary classification, multiclass classification, multivariate regression, and univariate time series) and three architectures (FFNN, CNN, and LSTM), it outperformed early stopping and the maximum-epochs criterion on every dataset, with improvements of up to 1.56% under oracle parameter selection where conventional criteria yielded negligible gains.

The method was shown to be applicable to a wide variety of supervised NN training problems, in which the validation loss correlates nontrivially with the training loss. Its architecture-agnostic formulation makes it a drop-in replacement for early stopping in many deep learning pipelines. The contribution is bounded by a three-parameter search space (window, threshold, and patience), larger than that of early stopping. Suitable settings are nonetheless reachable: Bayesian optimisation (BO) with gradient-boosted regression trees (GBRT) located them within 35 to 60 evaluations. The best window, threshold, and patience are task-dependent and are therefore set per problem rather than fixed to universal defaults, with the grid ranges used here serving only as starting points for that search. Natural extensions include applying the CDSC to more models and forecasting tasks such as PV generation and load forecasting. Integration with PyTorch and TensorFlow as callbacks would lower the adoption barrier for practitioners.

In *Thesis I.2* I developed a hyperparameter-optimised LSTM pipeline for HUPX DAM EPF. A grid search over lookback (6 h to 168 h), network width (6 to 168 neurons per layer), normalisation range, activation, and regularisation strategy identified an optimal configuration of 168 neurons, 48 h lookback, linear activation with $[0, 1]$ normalisation, and no weight-based regularisation, achieving MAE = 7.48 EUR/MWh. Weight-based regularisation degraded performance, motivating integration of the CDSC. Extended evaluation on the Germany (GER, EPEX SPOT), Singapore (SP, NEMS), and United Kingdom (UK,

N2EX) DAMs matched or outperformed early stopping in 10 of 12 configurations, with the advantage concentrated at the longer prediction horizon and the largest reduction, 16.7% in test mean absolute error (MAE), on SP at $n = 40$.

This formulation was shown to be valid for DAM EPF tasks at 30- and 60-minute resolution given several years of historical prices, the setting in which it was tested. The CDSC integration was shown to be beneficial for LSTM forecasting pipelines whose joint training and validation loss dynamics allow for correlation monitoring. The advantage was most reliable where the joint training and validation loss dynamics were informative and at the longer prediction horizon, where it was both largest and statistically significant. The principal caveat is that the optimal CDSC parameter configuration was specific to the DAM, since each DAM produced a different optimised CDSC threshold ranging from $\mu_{\text{cdsc}} = 0.0$ to $\mu_{\text{cdsc}} = 0.7$. Future research will focus on identifying cross-market insights across additional DAMs, and testing probabilistic forecasting model variants that propagate uncertainty into downstream trading decisions for risk-aware bidding.

In *Thesis I.3* I presented the first systematic context-length study of pre-trained Transformer foundation models for direct multi-step DAM forecasting. Six models (Chronos 2, three Chronos-Bolt variants, TimesFM 2.5, and Moirai-2) were evaluated on GER, SP, and UK at short ($n = 4$) and medium ($n = 40$) horizons, with L swept across 24 to 2048 timesteps in the zero-shot regime. The best-performing foundation model outperformed the best LSTM baseline in every scenario, with test MAE reductions of 40.1% and 30.0% on GER, 18.4% and 19.0% on SP, and 40.8% and 26.4% on UK (at $n = 4$ and $n = 40$); Chronos 2 was the strongest variant in five of the six scenarios. The foundation models also surpassed the conventional statistical, linear, and neural baselines, and the ranking held on a held-out 2025 robustness year on which no LSTM configuration matched the leading foundation models.

The framework is appropriate where forecasts must be produced quickly and without task-specific training across structurally distinct DAMs, provided the available history spans at least the longest tested context window. The performance hierarchy and the calendar-time scale of the optimal context length (roughly one week to a few months) are expected to carry over to neighbouring time-series forecasting tasks, though this was not tested here. The domain-independent ingredients are the calibration of context length in calendar time, the correlation-based training monitoring of the CDSC, and the direct multi-step formulation; the model ranking and the market-specific price dynamics would require re-tuning before transfer to other traded series such as commodities or foreign-exchange rates. Two assumptions constrain its use: zero-shot evaluation only, with no fine-tuning of any model on DAM data, and validation on three DAMs at two horizons. However, several of the tested foundation models support fine-tuning and in-context learning to refine results. A promising direction lies in domain-adaptive fine-tuning of the leading foundation models on DAM corpora, evaluation across additional European and global DAMs at intermediate horizons, and quantile or distributional outputs that preserve the zero-shot deployment advantage while supplying the uncertainty signals required for trading.

The second objective (Objective 2) was the development of optimisation tools for energy trading strategies for battery-enhanced renewable power plants, the stage that converts price forecasts into bids on the DAM. Two complementary algorithms addressed this objective, framed against a shared parametrised simplified DC-coupled hybrid inverter system model that supports both.

In *Thesis II.1* I proposed Optimised Energy Trading by Gradient Descent (OptEnT-GD), an adaptive gradient-descent-based method for 24-hour DAM energy trading optimisation, together with a parametrised DC-coupled hybrid inverter system model capturing battery capacity, storage efficiency η_{st}, inverter power P_{max}, and PV generation. OptEnT-GD identifies the most profitable transaction pair at each step and exploits the objective function's initial monotonicity through an adaptive step-size rule. Across 42 system configurations (three inverter powers, seven battery capacities, and two storage efficiencies), it achieved

up to 46.26% profit increase under predicted day-ahead prices on the configurations with 5% storage loss, with prediction error reducing the profit increase by only 1.93 to 6.93 percentage points relative to the oracle scenario, and the optimal minimum step size shifted from $\alpha_{\min} = 25$ (oracle) to $\alpha_{\min} = 50$ (predicted), evidencing implicit regularisation against prediction error (Energy 250, 2022, D1).

The method can be applied whenever the trading objective is continuous, differentiable, and exhibits initial monotonicity with respect to the chosen step size. It was tested with a simplified parametrised storage model with hard inverter and battery limits, but extensions to other models and objectives are possible. The main application targets are 24-hour granularity DAMs with hourly resolution and standard gate-closure timing. The contribution is primarily constrained by the initial monotonicity assumption, which may fail under non-convex tariff structures, nonlinear system models, coupled multi-day storage decisions, or once the objective is enriched with battery degradation, transaction costs, more detailed inverter losses, or other nonlinear operational penalties. In the profit objective evaluated here the assumption held, which is why OptEnT-GD converged reliably; it becomes binding only as these more realistic effects are introduced, where a derivative-free method serves as the fallback. Its other limitation is the simplified DC-coupled model, which utilises several assumptions that deviate from real-world operation to remain simple and fast. The main direction for future work is to improve the system model and validate it on a real-world system. After the model is validated, more complex tariff structures and objective functions can be developed that account for economic, engineering, and sustainability metrics, not just profit.

In *Thesis II.2* I proposed Optimised Energy Trading by Differential Evolution (OptEnT-DE), an adaptation of differential evolution [68] to the energy trading domain, together with a system sizing guide derived from the same 42-configuration evaluation. OptEnT-DE does not require the same monotonicity assumption as OptEnT-GD and offers a viable alternative when gradient information is unreliable. Comparative evaluation showed approximately 3% lower maximum profit than OptEnT-GD at the largest battery capacity and a 25–30% stochastic spread between the minimum and maximum profits across 100 runs. The sizing guide demonstrated that systems with larger battery-capacity-to-inverter-power ratios ($Q_{\max}/P_{\max}$) yield proportionally greater percentage profit increases, with maximum gains exceeding 50% compared to the immediate-dispatch strategy.

This formulation is valid wherever gradient information is unreliable, the objective is non-convex, or the monotonicity assumption underlying OptEnT-GD fails. The sizing guide applies to operators selecting between battery and inverter capacity at the procurement stage of a battery-enhanced renewable power plant. The principal caveat is the stochastic spread of 25% to 30% across 100 runs, which means a single OptEnT-DE execution provides no profit guarantee; the limitations from the simplified system model and 24-hour horizon inherited from OptEnT-GD apply equally here, and the sizing guide assumes the present DC-coupled architecture. A promising direction is reducing the negative effects of the algorithm's stochastic nature.

The third objective (Objective 3) was the physical design of PV and battery-enhanced energy systems under real-world operational, regulatory, and sustainability constraints, the stage on which the trading strategies of Objective 2 ultimately operate. Two methodologies addressed it: one for industrial multi-string installations under feed-in restrictions and one for semi-autonomous structures with tight physical constraints.

In *Thesis III.1* I proposed the dichotomous search-based string optimisation (DSSO), a consumption-aware string optimisation algorithm for multi-string PV energy systems under feed-in restrictions, together with an evaluation framework coupling simulated PV generation with year-long measured load at 15-minute resolution. DSSO bisects the feasible string-power interval $[P^{\text{str}}_{\min}, P^{\text{str}}_{\max}]$ using the sign of the objective's differential, converging in $O(\log_2((P^{\text{str}}_{\max} - P^{\text{str}}_{\min})/\varepsilon))$ iterations per string for any unimodal objective. Benchmarking on a Hungarian Small-to-Medium Enterprise (SME) dataset with four surface configurations

and two tariff structures yielded, under the fixed tariff, 14.8% to 17.5% lower E_{grid} than the E_{lost} limit method, approximately threefold higher μ_{sr}, and 12.9% to 15% lower CO_2 emissions at 37% to 45% of the Sum Yearly Consumption Matching (SYCM) system size; the DAM tariff returned smaller systems still.

The framework is appropriate wherever a multi-string PV installation faces feed-in restrictions and reliable load data is available at sub-hourly resolution. The unimodal-objective assumption was verified for NPV under both tariff structures tested; applying DSSO to other economic or sustainability metrics requires the same one-dimensional unimodality check first. The contribution is bounded by the unimodal-objective requirement, which excludes multi-modal objectives arising from complex tariff structures or non-convex incentive schemes. The comparison is scoped to standard engineering sizing rules (SYCM and the E_{lost} limit method); the claim is therefore that DSSO improves over these rules in the studied context, while competitiveness against optimisation-based sizing tools such as HOMER, PVSyst workflows, mixed-integer linear programming (MILP) formulations, or metaheuristic design is reserved for future work. Natural extensions include testing metaheuristic methods for handling non-unimodal objectives, broadening the set of metrics with economic, engineering, and sustainability-based metrics, and validating on a wider corpus of industrial load profiles.

In *Thesis III.2* I presented a comprehensive design methodology for semi-autonomous energy systems: the Stochastic Consumption Estimation Algorithm (SCEA) for year-long load-profile synthesis from appliance-level properties, a coupled thermal-electrical system model that accounts for the environmental heat exchange specific to structures with limited thermal mass, and a joint PV-storage optimisation that selects, among the configurations within a tolerance of the highest attainable self-reliance, the one leaving the least generation unutilised under the zero feed-in restriction. A simulation-based evaluation of 12 scenarios (four PV configurations and three storage sizes) on a real floating-home design, at 15-minute resolution over a full calendar year, identified the largest three-orientation PV array (26.65 kWp with a 20 kWh battery energy storage system, BESS) as reaching the highest self-reliance (91.43%) and an 86.2% grid-consumption reduction, but rejected it as oversized: relative to the available battery capacity, it left 22.14 MWh/year of generation unutilised under the feed-in restriction, so a smaller roof-only array was the more effective design.

The method applies whenever a building targets minimal grid reliance under tight weight, orientation, and thermal-mass constraints, with appliance-level usage data substituting for metered consumption. The thermal coupling is calibrated for water-surrounded structures but may generalise to lightweight prefabricated dwellings with similar heat-exchange profiles. The contribution is limited to simulated validation only, for a single semi-autonomous structure case study; physical validation of the constructed structure remains a future task. The thermal model decouples the thermal and electrical subsystems at the calibration boundary, and the SCEA stochastic profile is generated rather than learned from any metered baseline. Future research will focus on physical validation of the methodology on the completed semi-autonomous structure, end-to-end coupling of the thermal and electrical subsystems within the optimisation loop, and a metered-baseline comparison that anchors the SCEA profile against real consumption traces from comparable semi-autonomous structures.

Taken together, the methods were selected to match the structure of each subproblem rather than for uniformity, which offers a reusable selection guideline. The CDSC regularises NN training wherever the training and validation losses correlate informatively. For forecasting, pre-trained Transformer foundation models were chosen for their zero-shot accuracy across structurally distinct DAMs without per-market training, with the tuned LSTM as the in-domain baseline. For energy trading optimisation, OptEnT-GD exploits the initial monotonicity of the profit objective, and OptEnT-DE serves as the derivative-free fallback when that monotonicity cannot be assumed. For physical string sizing, dichotomous search was chosen for its logarithmic convergence and its physically interpretable single-panel

tolerance on objectives that are unimodal in string power. The common principle is to exploit problem structure where it exists (monotonicity, unimodality, and informative loss correlation) and to fall back on population-based or pre-trained general-purpose methods where it does not.

Returning to the motivation that opened this dissertation, the integration of weather-dependent renewables into the European grid demands accuracy at every link of the chain that converts intermittent generation into traded electricity. The seven contributions presented here strengthen each link in isolation, the long-term direction lies in closing the loop between them through coupled validation on a single DAM and site, so that forecasting, trading, and physical design optimisation reinforce one another in a unified decision-support workflow for prosumers and energy communities.

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] Directive (EU) 2023/2413 of the European Parliament and of the Council of 18 October 2023 amending Directive (EU) 2018/2001, Regulation (EU) 2018/1999 and Directive 98/70/EC as regards the promotion of energy from renewable sources, and repealing Council Directive (EU) 2015/652. *Official Journal of the European Union*, L 2023/2413, 31.10.2023, 2023. OJ L 2023/2413, 31.10.2023, CELEX: 32023L2413.

[3] International Energy Agency. Renewables 2024: Analysis and forecast to 2030. Technical report, IEA, Paris, 2024.

[4] Janina C. Ketterer. The impact of wind power generation on the electricity price in Germany. *Energy Economics*, 44:270–280, 2014.

[5] Johanna Cludius, Hauke Hermann, Felix Chr. Matthes, and Verena Graichen. The merit order effect of wind and photovoltaic electricity generation in Germany 2008–2016: Estimation and distributional implications. *Energy Economics*, 44:302–313, 2014.

[6] 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.

[7] 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.

[8] Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293:116983, 2021.

[9] Rafał Weron. *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*. John Wiley & Sons, Chichester, 2006.

[10] Tian Sheng Allan Loi and Gautam Jindal. Electricity market deregulation in Singapore—initial assessment of wholesale prices. *Energy Policy*, 127:1–10, 2019.

[11] George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.

[12] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

[13] Javier Contreras, Rosario Espinola, Francisco J. Nogales, and Antonio J. Conejo. ARIMA models to predict next-day electricity prices. *IEEE Transactions on Power Systems*, 18(3):1014–1020, 2003.

[14] Timothy D. Mount, Yumei Ning, and Xiaobin Cai. Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, 28(1):62–80, 2006.

[15] Reinaldo C. Garcia, Javier Contreras, Marco van Akkeren, and Joao Batista C. Garcia. A GARCH forecasting model to predict day-ahead electricity prices. *IEEE Transactions on Power Systems*, 20(2):867–874, 2005.

[16] Antonio J. Conejo, Miguel A. Plazas, Rosario Espinola, and Ana B. Molina. Day-ahead electricity price forecasting using the wavelet transform and ARIMA models. *IEEE Transactions on Power Systems*, 20(2):1035–1042, 2005.

[17] Kam Fong Chan, Philip Gray, and Bastiaan van Campen. Electricity price modelling with distributed lag structures. *Journal of Banking & Finance*, 32(12):2529–2539, 2008.

[18] Jinxing Che and Jianzhou Wang. Short-term electricity prices forecasting based on support vector regression and auto-regressive integrated moving average modeling. *Energy Conversion and Management*, 51(10):1911–1917, 2010.

[19] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[20] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

[21] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, volume 30, pages 3146–3154, 2017.

[22] Malte Lehna, Fabian Scheller, and Helmut Herwartz. Forecasting day-ahead electricity prices: A comparison of time series and neural network models taking external regressors into account. *Energy Economics*, 106:105742, 2022.

[23] Bartosz Uniejewski, Jakub Nowotarski, and Rafał Weron. Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies*, 9(8):621, 2016.

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

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

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

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

[28] 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.

[29] Souhaib Ben Taieb, Gianluca Bontempi, Amir F. Atiya, and Antti Sorjamaa. A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. *Expert Systems with Applications*, 39(8):7067–7083, 2012.

[30] Grzegorz Marcjasz, Jesus Lago, and Rafał Weron. Neural networks in day-ahead electricity price forecasting: Single vs. multiple outputs. *arXiv preprint arXiv:2008.08006*, 2020.

[31] 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.

[32] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11106–11115, 2021.

[33] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems*, volume 34, pages 22419–22430, 2021.

[34] Bryan Lim, Sercan Ö. Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, 2021.

[35] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *International Conference on Learning Representations (ICLR)*, 2023.

[36] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11121–11128, 2023.

[37] 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.

[38] Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, et al. Chronos-Bolt: A fast and accurate foundation model for time series forecasting. *arXiv preprint arXiv:2410.20081*,

2024.

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

[40] 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.

[41] 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.

[42] Leonard J. Tashman. Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting*, 16(4):437–450, 2000.

[43] Christoph Bergmeir and José M. Benítez. On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191:192–213, 2012.

[44] Vitor Cerqueira, Luis Torgo, and Igor Mozetič. Evaluating time series forecasting models: An empirical study on performance estimation methods. *Machine Learning*, 109(11):1997–2028, 2020.

[45] Marcos López de Prado. *Advances in Financial Machine Learning*. John Wiley & Sons, Hoboken, NJ, 2018.

[46] 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.

[47] Chaouki Ghenai, Adel Merabet, Tareq Salameh, and Erola Colon Pigem. Grid-tied and stand-alone hybrid solar power system for desalination plant. *Desalination*, 435:172–180, 2018.

[48] Pablo García, Juan P. Torreglosa, Luis M. Fernández, and Francisco Jurado. Improving long-term operation of power sources in off-grid hybrid systems based on renewable energy, hydrogen and battery. *Journal of Power Sources*, 265:149–159, 2014.

[49] V. Dash and P. Bajpai. Power management control strategy for a stand-alone solar photovoltaic-fuel cell-battery hybrid system. *Sustainable Energy Technologies and Assessments*, 9:68–80, 2015.

[50] Raffaele Cozzolino, Laura Tribioli, and Gino Bella. Power management of a hybrid renewable system for artificial islands: A case study. *Energy*, 106:774–789, 2016.

[51] L. W. Chong, Y. W. Wong, R. K. Rajkumar, and D. Isa. An optimal control strategy for standalone PV system with battery-supercapacitor hybrid energy storage system. *Journal of Power Sources*, 331:553–565, 2016.

[52] Zehra Yumurtaci. Role of energy management in hybrid renewable energy systems: case study-based analysis considering varying seasonal conditions. *Turkish Journal of Electrical Engineering & Computer Sciences*, 21(Sup. 2):2258–2271, 2013.

[53] Yoash Levron, Josep M. Guerrero, and Yuval Beck. Optimal power flow in microgrids with energy storage. *IEEE Transactions on Power Systems*, 28(3):3226–3234, 2013.

[54] Xiaomin Wu, Weihua Cao, Dianhong Wang, and Min Ding. A multi-objective optimization dispatch method for microgrid energy management considering the power loss of converters. *Energies*, 12(11):2160, 2019.

[55] H. R. Baghaee, M. Mirsalim, G. B. Gharehpetian, and H. A. Talebi. Reliability/cost-based multi-objective Pareto optimal design of stand-alone wind/PV/FC generation microgrid system. *Energy*, 115:1022–1041, 2016.

[56] H. R. Baghaee, M. Mirsalim, G. B. Gharehpetian, and H. A. Talebi. Multi-objective optimal power management and sizing of a reliable wind/PV microgrid with hydrogen energy storage using MOPSO. *Journal of Intelligent & Fuzzy Systems*, 32(3):1753–1773, 2017.

[57] Ali Kashefi Kaviani, Hamid Reza Baghaee, and Gholam Hossein Riahy. Optimal sizing of a stand-alone wind/photovoltaic generation unit using particle swarm optimization. *Simulation*, 85(2):89–99, 2009.

[58] H. R. Baghaee, M. Mirsalim, G. B. Gharehpetian, and H. A. Talebi. A decentralized power management and sliding mode control strategy for hybrid AC/DC microgrids including renewable energy resources. *IEEE Transactions on Industrial Informatics*, 14(6):2601–2609, 2018.

[59] Wei Gu, Jun Wang, Shuai Lu, Zhao Luo, and Chenyu Wu. An online optimal dispatch schedule for CCHP microgrids based on model predictive control. *IEEE Transactions on Smart Grid*, 8(5):2332–2342, 2016.

[60] Muhammad Khalid, Ricardo P. Aguilera, Andrey V. Savkin, and Vassilios G. Agelidis. A market-oriented wind power dispatch strategy using adaptive price thresholds and battery energy storage. *Wind Energy*, 21(4):242–254, 2018.

[61] Dongxiao Wang, Jing Qiu, Luke Reedman, Ke Meng, and Loi Lei Lai. Two-stage energy management for networked microgrids with high renewable penetration. *Applied Energy*, 226:39–48, 2018.

[62] M. Khalid and Andrey V. Savkin. A model predictive control approach to the problem of wind power smoothing with controlled battery storage. *Renewable Energy*, 35(7):1520–1526, 2010.

[63] M. Khalid and Andrey V. Savkin. An optimal operation of wind energy storage system for frequency control based on model predictive control. *Renewable Energy*, 48:127–132, 2012.

[64] Alessandra Parisio, Evangelos Rikos, and Luigi Glielmo. A model predictive control approach to microgrid operation optimization. *IEEE Transactions on Control Systems Technology*, 22(5):1813–1827, 2014.

[65] Wenzong Zhuo. Microgrid energy management strategy with battery energy storage system and approximate dynamic programming. In *2018 37th Chinese Control Conference (CCC)*. IEEE, 2018.

[66] Tao Chen, Xiao-Ping Zhang, Jianzhong Wang, Jianwei Li, Cong Wu, Mengqi Hu, and Haihong Bian. Battery management strategy using LSTM and Deep Deterministic Policy Gradient for personalized microgrid requirements. *Applied Energy*, 238:1500–1509, 2019.

[67] Azim Heydari, Meysam Majidi Nezhad, Elham Pirshayan, Davide Astiaso Garcia, Farshid Keynia, and Livio De Santoli. Short-term electricity price and load forecasting in isolated power grids based on composite neural network and gravitational search optimization algorithm. *Applied Energy*, 277:115503, 2020.

[68] 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.

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

[70] 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.

[71] G. Liu, M. G. Rasul, M. T. O. Amanullah, and M. M. K. Khan. Techno-economic simulation and optimization of residential grid-connected PV system for the Queensland climate. *Renewable Energy*, 45:146–155, 2012.

[72] Y. Fan and X. Xia. A multi-objective optimization model for building-integrated photovoltaic retrofit. *Energy and Buildings*, 141:1–12, 2017.

[73] S. Diaf, G. Notton, M. Belhamel, M. Haddadi, and A. Louche. Design and techno-economical optimization for hybrid PV/wind system under various meteorological conditions. *Applied Energy*, 85(10):968–987, 2008.

[74] G. Ning, L. Junnan, D. Yansong, Q. Zhifeng, J. Qingshan, G. Weihua, and D. Geert. BIM-based PV system optimization and deployment. *Energy and Buildings*, 150:13–22, 2017.

[75] M. Mehrabankhomartash, M. Rayati, A. Sheikhi, and A. M. Ranjbar. Practical battery

size optimization of a PV system by considering individual customer damage function. *Renewable and Sustainable Energy Reviews*, 67:36–50, 2017.

[76] A. R. Starke, J. M. Cardemil, R. A. Escobar, and S. Colle. Multi-objective optimization of hybrid CSP+PV system using genetic algorithm. *Energy*, 147:490–503, 2018.

[77] A. Al-Karaghouli and L. L. Kazmerski. Optimization and life-cycle cost of health clinic PV system for a rural area in southern Iraq using HOMER software. *Solar Energy*, 84(4):710–714, 2010.

[78] Yulianta Siregar, Y. Hutahuruk, and Suherman. Optimization design and simulating solar PV system using PVSyst software. In *2020 4th International Conference on Electrical, Telecommunication and Computer Engineering (ELTICOM)*. IEEE, 2020.

[79] Taehoon Hong, Myeonghwi Lee, Choongwan Koo, Kwangbok Jeong, and Jimin Kim. Development of a method for estimating the rooftop solar photovoltaic (PV) potential by analyzing the available rooftop area using Hillshade analysis. *Applied Energy*, 194:320–332, 2017.

[80] Behzad Lashkari, Yuxiang Chen, and Petr Musilek. Energy management for smart homes—state of the art. *Applied Sciences*, 9(17):3459, 2019.

[81] Daniel Akinyele, Ramesh Rayudu, and Nirmal Nair. Grid-independent renewable energy solutions for residential use: The case of an off-grid house in Wellington, New Zealand. In *2015 IEEE PES Asia-Pacific Power and Energy Engineering Conference (APPEEC)*, pages 1–6. IEEE, 2015.

[82] Jakub Gorzka, Artur Karczewski, Wojciech Litwin, Karolina Matej-Łukowicz, Nicole Nawrot, Lucyna Nyka, and Łukasz Piątek. Zero-emissions, off-grid, autonomous houseboat–a case study of selected locations in europe. *Polish Maritime Research*, 31(3):42–60, 2024.

[83] Jakub Grela, Aleksander Skała, Dominik Latoń, and Katarzyna Bańczyk. Enhancing energy autonomy in an e-houseboat: integration of renewable energy sources with hybrid energy storage systems. *Energies*, 18(5):1080, 2025.

[84] Fahmida Ishaque, M. S. Ahamed, and M. N. Hoque. Design and estimation of low cost floating house. *International Journal of Innovation and Applied Studies*, 7(1):49, 2014.

[85] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, Cambridge, MA, 2016.

[86] M. S. Iyer and R. R. Rhinehart. A novel method to stop neural network training. In *Proceedings of the 2000 American Control Conference*, volume 2, pages 929–933. IEEE, 2000.

[87] Xing-xing Wu and Jin-guo Liu. A new early stopping algorithm for improving neural network generalization. In *2009 Second International Conference on Intelligent Computation Technology and Automation*, pages 15–18. IEEE, 2009.

[88] J. T. Lalis, B. D. Gerardo, and Y. Byun. An adaptive stopping criterion for backpropagation learning in feedforward neural network. *International Journal of Multimedia and Ubiquitous Engineering*, 9(8):149–156, 2014.

[89] C. M. Ennett, M. Frize, and N. Scales. Evaluation of the logarithmic-sensitivity index as a neural network stopping criterion for rare outcomes. In *4th International IEEE EMBS Special Topic Conference on Information Technology Applications in Biomedicine*, pages 207–210, 2003.

[90] Minh Ha Nguyen, Hussein A. Abbass, and Robert I. McKay. Stopping criteria for ensemble of evolutionary artificial neural networks. *Applied Soft Computing*, 6(1):100–107, 2005.

[91] Jingbo Zhu, Huizhen Wang, Eduard Hovy, and Matthew Ma. Confidence-based stopping criteria for active learning for data annotation. *ACM Transactions on Speech and Language Processing*, 6(3):1–24, 2010.

[92] Michael Altschuler and Michael Bloodgood. Stopping active learning based on predicted change of F measure for text classification. In *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*, pages 47–54. IEEE, 2019.

[93] Hideaki Ishibashi and Hideitsu Hino. Stopping criterion for active learning based on deterministic generalization bounds. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 386–397. PMLR, 2020.

[94] Maren Mahsereci, Lukas Balles, Christoph Lassner, and Philipp Hennig. Early stopping without a validation set. *arXiv preprint arXiv:1703.09580*, 2017.

[95] Lu Peng, Shan Liu, Rui Liu, and Lin Wang. Effective long short-term memory with differential evolution algorithm for electricity price prediction. *Energy*, 162:1301–1314, 2018.

[96] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, volume 24, pages 2546–2554, 2011.

[97] Machine Learning Group, ULB. Credit card fraud detection. `https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud`, 2016. Accessed: 2022-10-30.

[98] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.

[99] Anastasia Makarova, Huibin Shen, Valerio Perrone, Aaron Klein, Jean Baptiste Faddoul, Andreas Krause, Matthias Seeger, and Cedric Archambeau. Automatic termination for hyperparameter optimization. In *International Conference on Automated Machine Learning*, pages 7/1–7/16. PMLR, 2022.

[100] Hideaki Ishibashi, Masayuki Karasuyama, Ichiro Takeuchi, and Hideitsu Hino. A stopping criterion for Bayesian optimization by the gap of expected minimum simple regrets. In *International Conference on Artificial Intelligence and Statistics*, pages 6463–6497. PMLR, 2023.

[101] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2623–2631. ACM, 2019.

[102] David Harvey, Stephen Leybourne, and Paul Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291, 1997.

[103] Peter R. Hansen, Asger Lunde, and James M. Nason. The model confidence set. *Econometrica*, 79(2):453–497, 2011.

[104] 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.

[105] Amazon Science. Chronos-2 model card. Hugging Face model card, 2025. Accessed: 2026-08-04.

[106] Google Research. TimesFM 2.5 model card. Hugging Face model card, 2025. Accessed: 2026-08-04.

[107] Salesforce AI Research. Moirai 2.0-R small model card. Hugging Face model card, 2025. Accessed: 2026-08-04.

[108] Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C. Maddix, Syama Rangapuram, David Salinas, Jasper Schulz, Lorenzo Stella, Ali Caner Türkmen, and Yuyang Wang. GluonTS: Probabilistic and neural time series modeling in Python. *Journal of Machine Learning Research*, 21(116):1–6, 2020.

[109] Boris N. Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. In *International Conference on Learning Representations (ICLR)*, 2020.

[110] T. Huld, R. Müller, and A. Gambardella. A new solar radiation database for estimating PV performance in Europe and Africa. *Solar Energy*, 86:1803–1815, 2012.

[111] ENTSO-E. Actual generation per production type, 2025. Accessed: 2025-02-27.

[112] S. Schlömer, T. Bruckner, L. Fulton, E. Hertwich, A. McKinnon, D. Perczyk, et al. Annex III: Technology-specific cost and performance parameters. In *Climate Change 2014: Mitigation of Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2014.

[113] United Nations Economic Commission for Europe. Carbon neutrality in the UNECE region: Integrated life-cycle assessment of electricity sources, 2022. Accessed: 2025-02-14.

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

[115] Jose Aguilar, Juan Giraldo, Manuela Zapata, Andrés Jaramillo, Luis Zuluaga, and Maria D R-Moreno. Autonomous cycle of data analysis tasks for scheduling the use of controllable load appliances using renewable energy. In *2021 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 1862–1867. IEEE, 2021.

[116] Christian Tipantuna and Xavier Hesselbach. IoT-enabled proposal for adaptive self-powered renewable energy management in home systems. *IEEE Access*, 9:64808–64827, 2021.

[117] DB Cahyono. Floating house platform stability test with a lightweight concrete cover made of expanded polystyrene system (EPS). In *IOP Conference Series: Earth and Environmental Science*, volume 1321, page 012044. IOP Publishing, 2024.

[118] Alexandra Simon Ivett, Miklós Boleraczki, and Gábor Gyurika István. Development of a highly automated construction and production planning framework for the production of floating houses. In *EXPRES 2024 — 14th International Symposium on Exploitation of Renewable Energy Sources and Efficiency*, pages 43–47, 2024.

[119] I. Ranaweera and O-M. Midtgård. Optimization of operational cost for a grid-supporting PV system with battery storage. *Renewable Energy*, 88:262–272, 2016.