



Short-Term Active Power Forecasting Under Uncertainty for Smart Distribution Grid Operation

Design of a Power Prediction model

Birendra Shahi

MSc in Electronic Systems, ES10-1055, 2026-06

Master's Thesis



Short-Term Active Power Forecasting Under Uncertainty for Smart Distribution Grid Operation

Design of a Power Prediction model

Birendra Shahi

MSc in Electronic Systems, ES19-1055, 2026-06

Master's Thesis



Copyright © Aalborg University 2015

Here you can write something about which tools and software you have used for typesetting the document, running simulations and creating figures. If you do not know what to write, either leave this page blank or have a look at the colophon in some of your books.



AALBORG UNIVERSITY
STUDENT REPORT

Electronics and IT
Aalborg University
<http://www.aau.dk>

Title:

Short-Term Active Power Forecasting Under Uncertainty for Smart Distribution Grid Operation

Theme:

Scientific Theme

Project Period:

Spring Semester 2026

Project Group:

ES10-1055

Participant(s):

Birendra Shahi

Supervisor(s):

Hans-Peter Schwefel
Rasmus Løvenstein Olsen

Copies: 1

Page Numbers: 86

Date of Completion:

4th June 2026

Abstract:

The growing penetration of rooftop solar photovoltaics in low-voltage residential distribution grids introduces bidirectional power flow and makes conventional load forecasting inadequate for congestion management. This thesis proposes a short-term active power forecasting pipeline combining Long Short-Term Memory (LSTM) networks with stochastic state estimation to deliver calibrated, uncertainty-aware grid operation support. Sixteen LSTM networks are trained on one year of substation measurements augmented with ERA5 weather features, achieving MAE = 18.95 kW and a 95 % confidence interval hit-rate of 99.5 % on the December 2023 test set, and HR = 96.1 % on January 2024 cross-year validation. The forecast is propagated through a Kirchhoff-consistent nine-node grid and passed to a constrained maximum likelihood state estimator, producing 95 % confidence ellipses with $HR_V = 95.0\%$ and $HR_I = 95.1\%$. A three-zone congestion risk classification demonstrates actionable decision support for distribution system operators managing LV grids with high distributed energy resource penetration.

List of Abbreviations

AC	Alternating Current
CCB	Customer Connection Box
DC	Direct Current
DG	Distributed Generation
DSO	Distribution System Operator
EV	Electric Vehicle
JB	Junction Box
LV	Low-Voltage
MV	Medium-Voltage
OGM	Observability Grid Model
PF	Power Factor
PV	Photovoltaic
XLPE	Cross-Linked Polyethylene (cable insulation)
DER	Distributed Energy Resources
BFS	Backward Forward Sweep
KCL	Kirchhoff's Current Law
KVL	Kirchhoff's Voltage Law
ACF	Autocorrelation Function
CI	Confidence Interval
MC	Monte Carlo
MLE	Maximum Likelihood Estimation
PDF	Probability Density Function
HR	Hit-Rate (fraction of true values inside the confidence interval)
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MBE	Mean Bias Error
MSE	Mean Squared Error
NMAE	Normalised Mean Absolute Error
RMSE	Root Mean Squared Error

AR AutoRegressive

ARIMA AutoRegressive Integrated Moving Average

LSTM Long Short-Term Memory

RNN Recurrent Neural Network

SARIMA Seasonal AutoRegressive Integrated Moving Average

EM Electricity Meter

GSE Grid State Estimation

API Application Programming Interface

CC Cloud Cover

ERA5 Fifth-Generation ECMWF Reanalysis (weather dataset)

SW Shortwave Radiation

Contents

List of Abbreviations	iv
Preface	x
1 Introduction	1
1.1 Background and Motivation	1
1.2 Problem Statement	2
1.3 Objectives	3
1.4 Related Work	4
1.5 Thesis Structure	5
2 Pre-Analysis	7
2.1 Smart Distribution Grids	7
2.1.1 Structure and Topology	7
2.1.2 Smart Metering and Measurement Model	8
2.1.3 Prosumer Grids and Bidirectional Power Flow	8
2.2 Active Power Forecasting	8
2.2.1 Problem Definition	8
2.2.2 Challenges at the Distribution Level	9
2.2.3 Why Deep Learning Over Classical Autoregression	9
2.3 Long Short-Term Memory Networks	9
2.3.1 Architecture	10
2.3.2 Dual-Output Design	11
2.4 Uncertainty Quantification	11
2.4.1 Two Sources of Uncertainty	12
2.4.2 Uncertainty Flow Through the Pipeline	12
2.4.3 Adaptive Confidence Interval	13
2.4.4 Estimator Output and Validation	14
2.5 Operational Applications	14

2.6	Toy Grid	15
2.6.1	Data Characteristics	15
2.7	Real Grid Dataset and Feature Correlation Study	17
2.7.1	Dataset Description	17
2.7.2	Annual Power Profile Analysis	17
2.7.3	Autocorrelation and Ramp Rate Analysis	18
2.7.4	Weather Data Acquisition	22
2.7.5	Feature Correlation Analysis and Selection	22
2.7.6	Dataset Split	25
3	Design and Methodology	26
3.1	Overall Pipeline Architecture	26
3.2	Module 1: LSTM Power Forecast Model	27
3.2.1	Design Philosophy — Two Networks Per Node	27
3.2.2	LSTM Architecture	29
3.2.3	Input Feature Construction	30
3.2.4	Solar Night Mask	30
3.2.5	Training and Validation Procedure	31
3.2.6	Confidence Interval Construction	31
3.3	Climatological Weather Proxy for Future Periods	32
3.3.1	The Problem	32
3.3.2	The Climatological Proxy Method	32
3.3.3	Implementation	32
3.4	Module 2: Grid Calculation	33
3.4.1	Purpose and Inputs	33
3.4.2	Synthetic Nine-Node Grid	33
3.4.3	Impedance Matrix Z	35
3.4.4	Constraint Matrix C from Kirchhoff's Laws	35
3.4.5	Voltage Phasors: Slack Bus and KVL Drops	36
3.4.6	Power Distribution to Nodes	37
3.4.7	Injection Current Derivation	37
3.4.8	KCL Upstream Summation	38
3.4.9	Noise Variance Computation	38
3.5	Module 3: Stochastic EM State Estimator	39
3.5.1	Overview and Purpose	39
3.5.2	Step 1 — EM Meter Model	40
3.5.3	Step 2 — Constrained MLE	40
3.5.4	Step 3 — Confidence Ellipses	41
3.5.5	Hit-Rate: Estimator Calibration	41

3.5.6	What uMeas and iMeas Are	41
3.5.7	What Was Changed in the Original Estimator Model’s Script	42
3.6	Module Integration and Data Files	43
3.6.1	node_forecast_output.mat	43
3.6.2	grid_inputs.mat	43
3.7	Chapter Summary	44
4	Results and Discussion	46
4.1	LSTM Training and Test-Set Performance	46
4.1.1	Training Configuration	46
4.1.2	December 2023 Test-Set Results	47
4.1.3	Baseline Comparison: AR Models vs LSTM	47
4.2	Validation Against 2024 Actual Data	49
4.2.1	Annual Validation Metrics	49
4.2.2	January 2024 — Detailed Analysis	49
4.2.3	Discussion of Annual Validation Pattern	51
4.2.4	Day-by-Day RMSE Assessment	53
4.2.5	RMSE-Based Current Uncertainty	54
4.2.6	External Validation Against True 2024 Power	54
4.3	Grid Calculation Results	56
4.3.1	Cable Current Magnitudes	56
4.3.2	Node Voltage Profile	56
4.4	State Estimation Results	58
4.4.1	Input to the State Estimator	58
4.4.2	Hit-Rate of Confidence Ellipses	58
4.5	Congestion Management Application	59
4.5.1	Two Directions of Congestion	59
4.5.2	Three-Zone Risk Classification	60
4.5.3	24-Hour Ahead Workflow	61
4.5.4	Case Study: December 15, 2023	61
4.6	Discussion	62
4.6.1	Model Strengths	62
4.6.2	Limitations and Their Origins	63
4.6.3	Comparison with Literature	63
5	Conclusion and Future Work	64
5.1	Summary of Contributions	64
5.2	Conclusions	66

5.3	Future Work	66
5.4	Closing Remarks	68
Bibliography		69
A PreAnalysis		72
A.1	MATLAB Pre-Analysis and Validation Scripts	72
A.1.1	Data Loading and Preprocessing	72
A.1.2	Autoregressive Power Forecasting	73
A.1.3	LSTM-Based Power Forecasting	73
A.1.4	Current Derivation and State Estimation	75
A.1.5	Results Evaluation and Visualisation	75
A.1.6	AR–LSTM Forecast Comparison	76
A.1.7	Discussion of Additional Forecast Comparison Results	77
A.1.8	Summary of the MATLAB Workflow	78
B Real Distribution Grid Dataset		79
B.1	Exploratory Data Analysis of the Real Distribution Grid Dataset	79
B.1.1	Monthly Consumption and Generation Patterns	79
B.1.2	Average Daily Load Profile	80
B.1.3	Seasonal Variations in Daily Operation	81
B.1.4	Load Duration Analysis	82
B.1.5	Reactive Power Characteristics	83
B.1.6	Consumption–Generation Operating Space	83
B.1.7	Weekly Behaviour	83
B.1.8	Monthly Daily Profiles	84
B.1.9	Statistical Power Distributions	85
B.1.10	Forecasting Motivation	85

Preface

This thesis was carried out during the final semester of the Master of Science (MSc) programme in Electronic Systems at Aalborg University. The work is situated within the fields of **Automation and Control** and **Engineering Systems**, and it aims to contribute to both the academic requirements of the programme and the investigation of optimal strategies for the configuration and operation of electrical distribution grids.

I would like to express our sincere gratitude to our supervisors, **Hans-Peter Schwefel** and **Rasmus L. Olsen**, for their continuous guidance, constructive feedback, and encouragement throughout the project. Their expertise and insightful discussions significantly shaped the direction and quality of this research. I also acknowledge their support in providing the datasets that were essential for the development and validation of the proposed model.



Birendra Shahi
<ZV11DE@student.aau.dk>

Chapter 1

Introduction

1.1 Background and Motivation

The ongoing transformation of the world’s energy supply is creating significant new challenges for electricity distribution networks. The rising share of rooftop solar photovoltaics (PV), heat pumps, and electric vehicles is shifting the classical concept of one-way, predictable power flow toward bidirectional flows at the distribution level [14, 17]. Low-voltage (LV) grids, originally designed to deliver electricity from a substation to passive consumers, now face bidirectional currents, voltage quality issues, and thermal overloading of cables [4] that were never dimensioned for reverse power flow.

Operating such a grid reliably requires two complementary capabilities:

- **Short-term active power forecasting:** predicting the net power at each grid node from one hour to a month ahead, so that the Distribution System Operator (DSO) can anticipate demand peaks and solar export surges before they cause cable overloads.
- **Probabilistic grid state estimation:** translating the forecasted power into physically consistent estimates of node voltages and cable currents, together with calibrated confidence regions that quantify estimation uncertainty.

Both capabilities are addressed in this thesis through a three-module sequential pipeline. First, a Long Short-Term Memory (LSTM) network forecasts the consumed and generated power at each grid node using ERA5

weather data — temperature, cloud cover, and shortwave radiation — as inputs. Second, the forecasted power is converted to complex current phasors at each node using the three-phase power formula and Kirchhoff's Current Law (KCL), propagated from leaf nodes to the substation root. The substation (N1) is treated as the slack bus at 400 V; all other nodes receive slightly lower voltages due to KVL cable voltage drops. Third, the derived phasors are passed to the stochastic state estimator of Olsen et al. [18], which adds realistic electricity-meter noise, solves a constrained maximum likelihood problem, and produces 95 % confidence ellipses for every node phasor.

The practical motivation for the pipeline is **proactive congestion management**. Congestion in an LV grid means a cable current exceeds its thermal rating (200 A for 95 mm² XLPE cables [11]), and it can occur in two directions: forward (high demand) and reverse (high solar export). The maximum recorded export in this dataset is 217.6 kW, which would produce a substation current of 314 A — 57 % above the rated limit. A DSO cannot react to congestion in real time; they must predict it and act preventively. The 60-minute ahead LSTM forecast, combined with a calibrated confidence interval validated at HR = 96.1 %, provides the quantified uncertainty that drives the SAFE / WARNING / RISK classification used for preventive action.

1.2 Problem Statement

The central challenge of this thesis is to design and validate a short-term active power forecasting pipeline that is both accurate enough for practical use and uncertainty-aware enough to support risk-based congestion management in a smart LV distribution grid.

The pipeline operates under the following constraints:

- Only substation-level power measurements are available for training (one year, 15-minute resolution, gridElementId 10, Neuburg, Bayern). Individual node measurements do not exist; node-level power is approximated by proportional distribution using fixed weights.
- The target forecast period may be a future month for which no weather data yet exists; a climatological proxy (same month of 2023) is used as the weather input.
- The state estimator is provided by the supervisor and must be integrated without modifying its estimation logic.

This leads to the following problem statement:

How can a weather-driven per-node LSTM forecasting model, combined with a KCL-based current derivation and a stochastic EM state estimator, provide calibrated uncertainty quantification sufficient to support proactive congestion management in a low-voltage distribution grid, when only substation-level historical data is available for training?

1.3 Objectives

The specific objectives of this thesis are:

1. **Data analysis and feature selection:** Perform a pre-analysis of the one-year substation power dataset (35,031 steps, January–December 2023) to characterise seasonal consumption and generation patterns, auto-correlation structure, and the correlation between power and ERA5 weather variables. Select the three most informative weather features for the LSTM input.
2. **Model comparison:** Implement an autoregressive (AR) baseline and a from-scratch LSTM model on the 2023 dataset and compare their 60-minute ahead point forecast accuracy, establishing that LSTM provides a calibrated confidence interval whereas AR does not.
3. **Per-node LSTM design:** Split the 2023 substation data to eight load nodes using fixed proportional weights, train 16 independent LSTM networks (two per node: consumed and generated power), and validate the per-node design on the December 2023 test set and the January 2024 cross-year holdout.
4. **Grid calculation:** Implement a KCL-based current derivation for a synthetic nine-node radial LV grid, computing per-node injection currents from the LSTM forecast power and propagating them from leaves

to the substation root. Apply KVL from the slack bus ($N1 = 400\text{ V}$) to compute the correct voltage drop at each downstream node.

5. **State estimator integration:** Integrate the forecast-derived phasors as the operating-point input to the supervisor’s stochastic EM state estimator, validate the 95% confidence ellipses through hit-rate analysis (target $HR \approx 95\%$), and clarify the distinction between the forecasted (measured) input and the true physical grid state.
6. **Congestion management framework:** Develop a 24-hour ahead congestion assessment framework using the RMSE-based current uncertainty ($\Delta I = 28.5\text{ A}$), classify each timestep into SAFE / WARNING / RISK zones, and demonstrate the framework on a case study day, distinguishing forward (demand) and reverse (export) congestion.

1.4 Related Work

Grid state estimation: State estimation in distribution networks has a long history, from classical weighted least squares methods [1] to more recent approaches tailored to LV grids with high DER penetration [5, 16]. The work of Olsen et al. [18] is directly used in this thesis: it develops a stochastic EM model that produces confidence ellipses for estimated node phasors using constrained maximum likelihood estimation under Gaussian meter noise, validated on a real European 400 V LV grid with 98 customers.

Short-term load forecasting: LSTM networks [8] have achieved strong performance on seasonal energy sequences by learning long-range dependencies through gated memory cells. Weather variables — particularly solar irradiance and ambient temperature — are widely recognised as the most informative external features for distribution-level forecasting [7, 20]. The use of separate models for consumed and generated power, rather than a single net-power model, is motivated by two complementary observations from the pre-analysis (Chapter 2). First, the two quantities follow **opposing seasonal patterns:** mean monthly consumption peaks at 67.7 kW in December and falls to 19.7 kW in June, while mean monthly generation peaks at 33.6 kW in June and drops to near zero in winter. A single net-power model trained on their difference would need to learn both patterns simultaneously from the same weather inputs, creating a learning

conflict that degrades accuracy for both components. Second, the quantities respond to weather in fundamentally different ways: consumption is negatively correlated with temperature (cold weather increases heating demand), while generation is positively correlated with shortwave solar radiation (clear skies drive PV output). Both claims are quantified in the correlation analysis of Section 2.7.5.

Congestion management in LV grids: As DER penetration increases, voltage and thermal constraints become binding more frequently and in both current directions [17]. Proactive congestion management — issuing preventive demand response or curtailment signals before a violation occurs — requires not only a point forecast but also a quantified uncertainty bound around it. The distinction between forward-flow congestion (high demand) and reverse-flow congestion (high solar export) is operationally important: the two cases require different DSO interventions (load shifting vs. PV curtailment) and have different risk profiles given the asymmetry between maximum recorded consumption (147.2 kW) and maximum recorded generation (217.6 kW) in the dataset used.

1.5 Thesis Structure

The remainder of this thesis is organised as follows.

Chapter 2 presents the pre-analysis of the 2023 substation power dataset and ERA5 weather data. It characterises seasonal patterns, autocorrelation structure, and the correlation between power and weather variables. It compares an autoregressive baseline with the LSTM model and justifies the selection of temperature, cloud cover, and shortwave radiation as the three LSTM input features.

Chapter 3 describes the complete design of the three-module pipeline: the per-node LSTM forecast model with its training procedure and confidence interval construction; the nine-node synthetic grid calculation including the slack-bus voltage assumption, KVL voltage drops, and KCL current propagation; and the integration with the supervisor’s estimator model.

Chapter 4 reports all results: LSTM forecast accuracy on the December 2023 test set and January 2024 cross-year validation, comparison with AR baselines, node voltage profiles, state estimator hit-rates ($HR_V = 95.0\%$, $HR_I = 95.1\%$), and the 24-hour congestion management case study for December 15, 2023.

Chapter 5 summarises the contributions, draws conclusions from the completed work, and outlines directions for future work — in particular, the replacement of proportional weight-based node data with real per-node smart meter measurements when those become available.

Chapter 2

Pre-Analysis

This chapter provides the theoretical foundations and dataset characterisation required to contextualise the forecasting framework. Section 2.1 describes smart distribution grid structure. Section 2.2 introduces the active power forecasting problem. Section 2.3 reviews the LSTM architecture and motivates its selection. Section 2.4 examines uncertainty quantification. Section 2.5 discusses operational applications. Section 2.6 describes the OGM toy-grid dataset used for pipeline development. Section 2.7 presents the one-year real grid dataset, annual data analysis, ERA5 weather acquisition, and the feature correlation study that drives the feature selection documented in Chapter 3.

2.1 Smart Distribution Grids

2.1.1 Structure and Topology

A low-voltage (LV) distribution grid is the final stage of the electrical power system, connecting the secondary substation transformer to individual customers at 400 V line-to-line. The network is typically structured as a radial tree: power flows from a single root node through junction boxes (JBs) to customer connection boxes (CCBs) at the leaves. The electrical state is fully described by the complex voltage phasor $V_k = u_k e^{j\theta_k}$ and current phasor I_k at each node k , satisfying the Kirchhoff constraint:

$$\mathbf{C}\mathbf{x} = \mathbf{c} \quad (2.1)$$

where $\mathbf{x} \in \mathbb{C}^{2N}$ is the full phasor state vector, \mathbf{C} is the topology-and-impedance constraint matrix built from the impedance matrix \mathbf{Z} , and \mathbf{c}

encodes the operating point [18].

2.1.2 Smart Metering and Measurement Model

Smart meters record voltage magnitude u_k , current magnitude i_k , and local phase angle ϕ_k at 15-minute intervals. Each measurement contains Gaussian noise:

$$\tilde{u}_k = u_k + \varepsilon_{u,k}, \quad \varepsilon_{u,k} \sim \mathcal{N}(0, \sigma_{u,k}^2) \quad (2.2)$$

$$\tilde{i}_k = i_k + \varepsilon_{i,k}, \quad \varepsilon_{i,k} \sim \mathcal{N}(0, \sigma_{i,k}^2) \quad (2.3)$$

$$\tilde{\phi}_k = \phi_k + \varepsilon_{\phi,k}, \quad \varepsilon_{\phi,k} \sim \mathcal{N}(0, \sigma_{\phi}^2) \quad (2.4)$$

Accuracy parameters from Olsen et al. [18]: voltage relative error $\varepsilon_U = 1\%$, current relative error $\varepsilon_I = 3\%$, phase-angle noise $\sigma_{\phi} = 0.01$ rad. These specifications are used directly in Module 2 of the pipeline (Section 3.4).

2.1.3 Prosumer Grids and Bidirectional Power Flow

A *prosumer* simultaneously consumes and generates power. The net active power at node k :

$$P_{\text{net}}(t, k) = P_{\text{con}}(t, k) - P_{\text{gen}}(t, k) \quad (2.5)$$

is positive during net consumption (grid import) and negative during solar export (grid export). High PV penetration creates a characteristic daily inversion: overnight consumption, midday solar-export trough, and evening demand peak. In the 2024 dataset used for validation, this export pattern is present in 21.4% of all 15-minute intervals.

2.2 Active Power Forecasting

2.2.1 Problem Definition

Short-term active power forecasting predicts $P_{\text{net}}(t+h | t)$ at horizon h given all observations up to time t . At 15-minute resolution, $h = 4$ corresponds to 60 minutes ahead. The system must provide both a point forecast $\hat{P}(t+h)$ and a calibrated 95% prediction interval.

2.2.2 Challenges at the Distribution Level

LV forecasting is harder than transmission-level forecasting because of:

- **High variability** of individual customer loads;
- **Non-stationarity** introduced by new PV and EV connections between training and test periods; and
- **Strong weather dependence** of solar generation, requiring meteorological inputs that are not needed for pure demand forecasting.

2.2.3 Why Deep Learning Over Classical Autoregression

Classical AR(p) models express the current value as a linear combination of p past observations. Their analytical tractability is appealing, but they cannot capture:

- The nonlinear interaction between temperature and heating demand [19];
- The multiplicative relationship between solar radiation, cloud cover, and PV generation [19]; or
- The regime change between winter (heating-dominated import) and summer (solar-dominated export) [13].

LSTM networks address all three limitations [12, 13, 19]. The autocorrelation analysis in Section 2.7.3 confirms multi-scale periodicity ($\rho_{96} = 0.814$, $\rho_{672} = 0.763$) that motivates a model with long-range memory and exogenous weather inputs. LSTM networks naturally satisfy both requirements [12].

2.3 Long Short-Term Memory Networks

Before describing the LSTM architecture in detail, Figure 2.1 gives an overview of the complete system and shows where each component sits in relation to the others.

The pipeline has three sequential modules: Module 1 (LSTM, Section 3.2) forecasts the net power at each of the eight load nodes from ERA5 weather inputs. Module 2 (grid calculation, Section 3.4) converts the power forecast

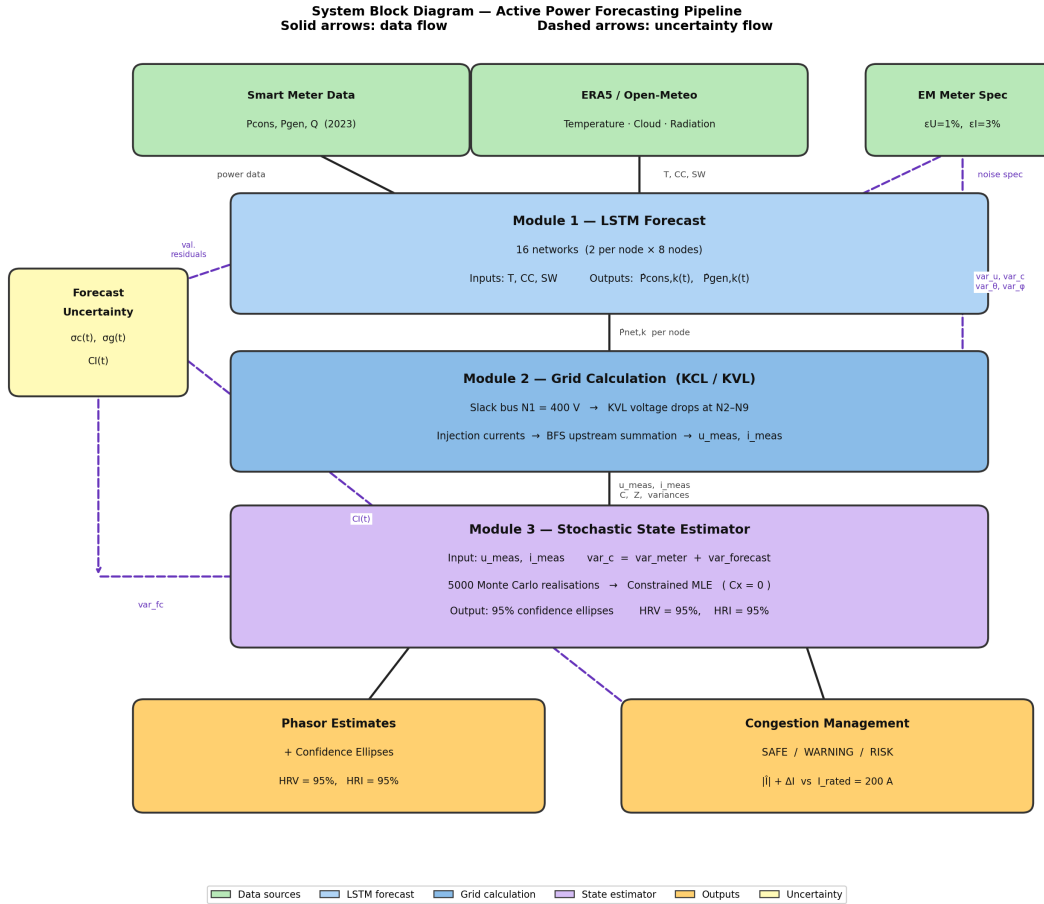


Figure 2.1: System block diagram of the active power forecasting pipeline

to voltage and current phasors using KCL and KVL on the nine-node synthetic grid. Module 3 (state estimator, Section 3.5) adds realistic EM meter noise, solves the constrained MLE for 5000 Monte Carlo realisations, and produces 95% confidence ellipses. The uncertainty block shows how the LSTM forecast residuals (σ_c, σ_g) feed both the adaptive confidence interval for congestion management and the current variance passed to the state estimator.

2.3.1 Architecture

LSTM networks address the vanishing gradient problem of vanilla RNNs through a gated memory architecture [8]. The complete LSTM update at

timestep t :

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (2.6)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (2.7)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (2.8)$$

$$\mathbf{g}_t = \tanh(\mathbf{W}_g \mathbf{x}_t + \mathbf{U}_g \mathbf{h}_{t-1} + \mathbf{b}_g) \quad (2.9)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t \quad (2.10)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (2.11)$$

The additive cell-state update enables learning of dependencies spanning hundreds of timesteps, making LSTMs well-suited to power time series with daily (96 steps) and seasonal periodicity.

2.3.2 Dual-Output Design

Two independent LSTM networks are trained: one for consumed power \hat{P}_{cons} and one for generated power \hat{P}_{gen} . This separation exploits the distinct physical drivers of each quantity:

- Consumption is primarily driven by temperature and time-of-day patterns (heating and occupancy);
- Generation is primarily driven by solar radiation (shortwave irradiance) and cloud cover.

A joint single-output model predicting net power directly would conflate these two mechanisms, increasing model complexity without gain. The feature correlation analysis in Section 2.7.5 quantifies these relationships and confirms the separation is appropriate.

2.4 Uncertainty Quantification

Accurate forecasting alone is insufficient for grid operation — the DSO also needs to know *how wrong* the forecast could be. This section describes the two distinct sources of uncertainty in the pipeline and how each is quantified and used.

2.4.1 Two Sources of Uncertainty

The pipeline contains two independent sources of uncertainty that must not be confused:

1. **Forecast uncertainty** — arises from the LSTM prediction error. Following the standard practice for prediction interval construction [9], it is quantified from the residuals on the November 2023 validation set \mathcal{V} :

$$\sigma_c = \text{std}\{\hat{P}_{\text{cons}}(t) - P_{\text{cons}}(t)\}_{t \in \mathcal{V}}, \quad \sigma_g = \text{std}\{\hat{P}_{\text{gen}}(t) - P_{\text{gen}}(t)\}_{t \in \mathcal{V}} \quad (2.12)$$

Substituting the trained model values ($\sigma_c = 13.164$ kW, $\sigma_g = 14.762$ kW) and horizon $h = 4$ gives the 95% prediction interval [9]:

$$\hat{P}_{\text{net}}(t+h) \pm 1.96 (\sigma_c + \sigma_g) \sqrt{h} = \pm 109.47 \text{ kW} \quad (2.13)$$

The \sqrt{h} factor accounts for accumulating uncertainty over the multi-step horizon. The forecast uncertainty is used for two purposes: it directly provides the congestion management CI, and it is converted to a current variance that feeds into the state estimator (Section 2.4.2).

2. **EM meter noise** — arises from the finite accuracy of smart electricity meters as specified in the EM model of Olsen et al. [18]. The meter accuracy is expressed as a percentage of the measured value at the 99% confidence level ($\varepsilon_U = 1\%$ for voltage, $\varepsilon_I = 3\%$ for current). Converting to standard deviations using the Gaussian quantile $r_0 = \Phi^{-1}(0.995) = 2.576$ [18]:

$$\sigma_U = \frac{V_{\text{nom}} \varepsilon_U}{r_0} = \frac{400 \times 0.01}{2.576} = 1.55 \text{ V}, \quad \sigma_I(t) = \frac{|\bar{I}|(t) \varepsilon_I}{r_0} \quad (2.14)$$

The EM meter noise is added inside the state estimator by `meterModel_1` and drives the size of the confidence ellipses.

2.4.2 Uncertainty Flow Through the Pipeline

Figure 2.2 shows how the two uncertainty sources propagate through the pipeline.

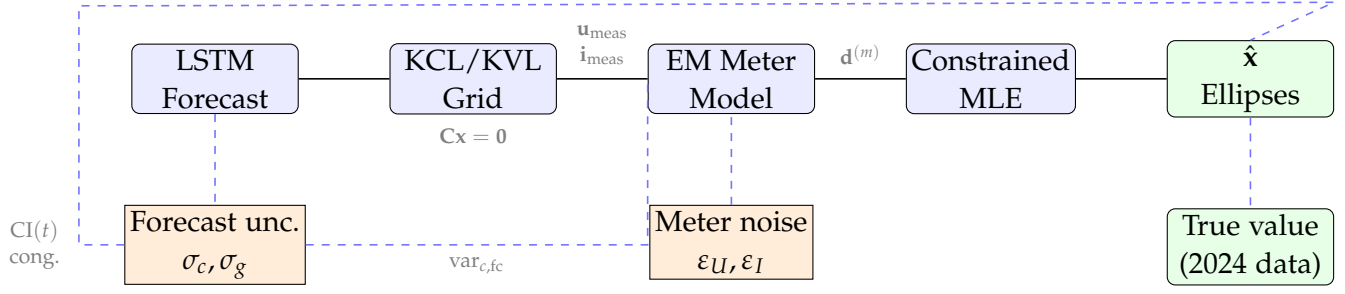


Figure 2.2: Uncertainty flow through the pipeline. Solid arrows: data flow. Dashed arrows: uncertainty flow. Forecast uncertainty (σ_c, σ_g) feeds both the congestion CI and the estimator variance $\text{var}_{c,fc}$. EM meter noise drives the ellipse size. Output \hat{x} is compared against 2024 actual data ($\text{HR}_V = 95\%$, $\text{HR}_I = 95\%$).

The total current variance passed to `estimateGridState` combines both sources [18]:

$$\text{var}_{I,k}(t) = \underbrace{\left(\frac{|\bar{I}|(t) \varepsilon_I}{r_0} \right)^2}_{\text{EM meter noise}} + \underbrace{\left(\frac{w_k \sqrt{\sigma_c^2 + \sigma_g^2} \times 1000}{\sqrt{3} V_{\text{nom}}} \right)^2}_{\text{Forecast uncertainty}} \quad (2.15)$$

where w_k is the power allocation weight for node k . The forecast component (28.5 A at the substation) dominates the meter noise (0.58 A at typical loading) by a factor of ≈ 49 , so the confidence ellipses primarily reflect LSTM prediction error rather than instrument accuracy.

2.4.3 Adaptive Confidence Interval

The fixed CI of ± 109.47 kW treats all hours and seasons equally. In practice the LSTM is more uncertain during summer solar events than during stable winter nights. Following the conditional prediction interval approach [9], an adaptive CI is computed by partitioning the November 2023 residuals by **time-of-day slot** (0–95, one per 15 min) and **season** (winter / spring / summer / autumn):

$$\text{CI}(t) = 1.96 \left(\sigma_c(\text{slot}(t), \text{season}(t)) + \sigma_g(\text{slot}(t), \text{season}(t)) \right) \sqrt{h} \quad (2.16)$$

This replaces one fixed value with a $4 \times 96 = 384$ -element lookup table stored in `adaptive_sigma.mat`. The CI narrows to $\approx \pm 20$ kW during quiet winter nights and widens to $\approx \pm 210$ kW during peak summer solar hours.

2.4.4 Estimator Output and Validation

The state estimator solves the constrained MLE problem [18]:

$$\hat{\mathbf{x}}^{(m)} = \arg \min_{\mathbf{x}} \|\mathbf{d}^{(m)} - \mathbf{x}\|_{\Sigma^{-1}}^2 \quad \text{s.t.} \quad \mathbf{C}\mathbf{x} = \mathbf{0} \quad (2.17)$$

for each of the 5000 Monte Carlo realisations $\mathbf{d}^{(m)}$. The confidence ellipses are sized using the chi-squared quantile [18]:

$$\alpha_{\text{factor}} = \chi_{0.95}^2(2) = 5.9915 \quad (2.18)$$

ensuring 95 % theoretical coverage in the 2D phasor plane.

The estimator output is validated on two complementary levels:

- **Internal calibration** (Monte Carlo hit-rate): do 95 % of ellipses contain the forecast-derived operating point? Result: $\text{HR}_V = 95.0\%$, $\text{HR}_I = 95.1\%$ — matching the paper base case [18].
- **External accuracy** (actual 2024 measurements): how close is the forecast-derived current to reality? Result: current $\text{MAE} \approx 37.6 \text{ A}$, $\text{MBE} \approx +1.8 \text{ A}$ (near-zero systematic bias).

Calibration confirms the ellipses are correctly sized for the specified meter noise; external accuracy confirms the forecast-derived input has negligible systematic error. Both metrics together provide a complete uncertainty assessment.

2.5 Operational Applications

Probabilistic power forecasts directly support three DSO use-cases:

Congestion management. The upper confidence bound $\hat{P} + \text{CI}$ is translated to a cable current via $|I| = |S|/(\sqrt{3} V_{\text{nom}})$ and compared against the thermal rating I_{rated} . If the upper bound exceeds the limit, a risk warning is raised before congestion occurs.

Demand response scheduling. The lower confidence bound quantifies the minimum expected solar export, allowing the DSO to safely schedule controllable loads to absorb surplus generation without risking reverse-flow overloads.

State estimation input. The point forecast provides the injection currents that seed the stochastic state estimator (Module 3 of the pipeline), which propagates forecast uncertainty through the Kirchhoff equations to produce 95% confidence ellipses at all grid nodes.

2.6 Toy Grid

Initially, the forecast model is trained with the 24 hours data from the previous semester’s project [2]. The project work is carried out with some real grid dataset (from the previous semester) so that some idea can be acquired about the future of the project.

2.6.1 Data Characteristics

The OGM data covers a single 24-hour day, yielding 96 timesteps per node. Table 2.1 summarises the substation-level statistics.

Table 2.1: OGM toy-grid summary statistics — Substation Node 91, 1 August 2024.

Quantity	Value	Unit
Timesteps	96 (15-min intervals)	
Net power range	−196.4 to +70.0	kW
Peak solar export	−196.4	kW (at 12:30)
Peak evening demand	+70.0	kW (at 18:30)
Individual node range	−3.8 to +1.1	kW

Forecast Model Evaluation on the OGM Dataset

Figure 2.3 shows the AR(2) and LSTM forecast results for Substation Node 91 trained and evaluated on the single-day OGM dataset, with the train/test split at hour 16 (dashed green line).

The results clearly illustrate the limitations of training on a single day of data. The AR(2) model achieves a mean absolute error (MAE) of 14.7 kW, substantially outperforming the LSTM, which yields an MAE of 26.8 kW — nearly twice the error of the classical baseline. Notably, the LSTM struggles to track the sharp evening demand transition following the solar export

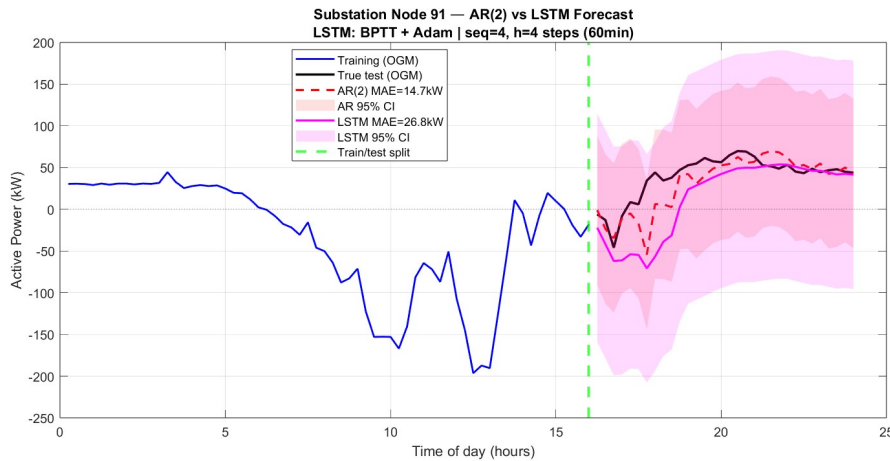


Figure 2.3: AR(2) vs. LSTM forecast for Substation Node 91 on the OGM toy-grid dataset.

period, and its 95% confidence interval widens dramatically over the test horizon, indicating high predictive uncertainty.

This performance gap is unsurprising given the severe data constraints. With only 96 timesteps available and a train/test split at timestep 64, the LSTM is trained on approximately 60 samples — far fewer than the tens of thousands typically required to tune its 18,753 parameters reliably. The model consequently overfits the limited training signal and fails to generalise to the test period.

The single-day OGM dataset is sufficient to:

- Configure the 122-node grid state estimator (static topology and one operating point); and
- Validate the complete forecast-to-current-to-state pipeline end-to-end with a realistic European grid.

It is *not* sufficient for training a reliable LSTM model: 96 training samples against 18,753 LSTM parameters is far below the recommended minimum. This motivates the use of the real 2023 substation dataset for all LSTM training (Section 2.7).

2.7 Real Grid Dataset and Feature Correlation Study

2.7.1 Dataset Description

The primary dataset is a one-year OGM export from a European DSO covering calendar year 2023. The file `power_export_gridElementId_10.csv` records the aggregated power exchange at Grid Element 10 (substation level) at 15-minute intervals, yielding 35,031 timesteps. This is distinct from the OGM toy-grid node files: it provides substation-aggregated consumption and generation totals, not individual node phasors.

Table 2.2: Real grid dataset summary statistics — Grid Element 10, 2023.

Variable	Mean	Max	Min	Unit
Consumed (P_{con})	37.4	147.2	0.0	kW
Generated (P_{gen})	16.5	217.6	0.0	kW
Net power (P_{net})	20.9	147.2	-96.2	kW
Inductive reactive power	9.6	19.6	0.0	kVAr
Capacitive reactive power	0.4	20.0	0.0	kVAr
Total consumed	327.6 MWh/year			
Total generated	144.5 MWh/year			

2.7.2 Annual Power Profile Analysis

Figure 2.4 shows the daily-averaged consumed and generated power across 2023. Consumption peaks in winter (mean 57–68 kW in December–January, driven by electric space heating) and falls in summer (mean 19–24 kW in June–July). Generation follows the complementary solar pattern, rising from near-zero in December to a mean of 28–34 kW in June–July.

Table 2.3 presents the monthly mean consumed, generated, and net power for 2023. From May to September the grid operates in net export mode (negative mean net power), as solar generation exceeds consumption. This seasonal inversion is the primary modelling challenge and confirms the need for weather-informed features rather than a purely time-series approach.

Figures 2.5 and 2.6 show the full year as heatmaps of hour-of-day (rows) versus day-of-year (columns). The consumed power heatmap reveals persistent morning and evening demand peaks throughout the year. The gen-

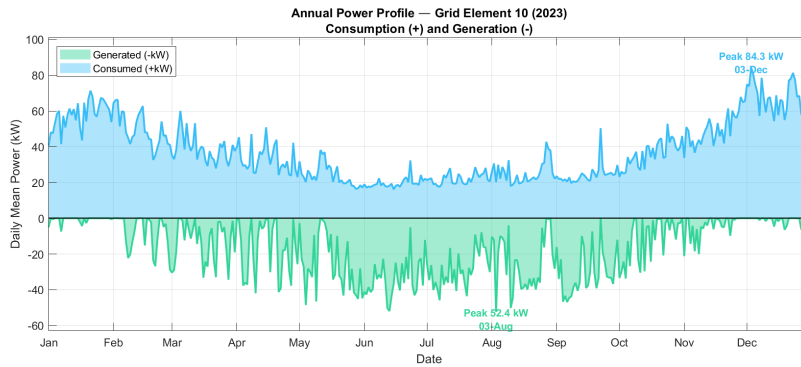


Figure 2.4: Annual daily-averaged power profile — Grid Element 10, 2023

Table 2.3: Monthly mean power statistics — Grid Element 10, 2023. Negative net power indicates net solar export.

Month	Consumed (kW)	Generated (kW)	Net (kW)
January	57.7	0.7	+57.0
February	49.8	7.5	+42.3
March	38.9	13.0	+25.9
April	33.3	16.6	+16.7
May	24.0	25.4	-1.5
June	19.7	33.6	-13.9
July	22.1	28.8	-6.7
August	25.5	25.1	+0.4
September	24.4	30.2	-5.8
October	36.2	11.4	+24.9
November	50.1	4.4	+45.8
December	67.7	1.0	+66.7

erated power heatmap confirms that solar generation is concentrated in midday hours of spring and summer and near-zero throughout winter — the key temporal pattern that the LSTM must learn.

2.7.3 Autocorrelation and Ramp Rate Analysis

Autocorrelation Structure

Figure 2.7 shows the autocorrelation function (ACF) of consumed power at lags up to one week. Key values:




Figure 2.5: Consumed power: morning and evening peaks throughout the year.




Figure 2.6: Generated power: midday solar peak, dominant May–September.

- $\rho(1) = 0.963$ (15 min) — very strong short-term persistence: the power in the next 15 minutes is almost perfectly predicted by the current value alone.
- $\rho(4) = 0.867$ (1 hour) — correlation remains high at 1-hour lag, confirming that load follows smooth intraday ramps.
- $\rho(32) = 0.148$ (8 hours) — correlation has dropped substantially at the 8-hour horizon, meaning that an 8-hour lookback window captures the most recent smooth ramp history without carrying stale information from 12 hours earlier.
- $\rho(96) = 0.814$ (1 day) — a strong daily seasonality spike: consumption today is highly correlated with consumption at the same time yesterday.
- $\rho(672) = 0.763$ (1 week) — a weaker but still significant weekly seasonality.




Figure 2.7: Autocorrelation function of consumed power. Left: short lags (0–12 h). Right: full week (0–672 steps). Dashed red lines: 95% confidence bounds.

Justification for the 8-hour lookback window ($L = 32$). The ACF shows two distinct regimes. For lags up to approximately 16 steps (4 hours), the correlation decays smoothly from $\rho(1) = 0.963$ to $\rho(16) = 0.447$, reflecting the intraday ramp dynamics (heating onset, occupancy changes). Beyond 16 steps the correlation drops further to $\rho(32) = 0.148$ before rising again at the daily peak ($\rho(96) = 0.814$).

A lookback window of $L = 32$ steps (8 hours) is chosen because it covers the full smooth ramp regime where autocorrelation is meaningful for short-term prediction, without extending into the stale mid-day region ($\rho \approx 0.14$) that adds noise rather than signal. The dominant daily seasonality ($\rho_{96} = 0.814$) is **not** captured directly by the lookback window, since 32 steps cannot reach 24 hours back. Instead, this daily cycle is captured **implicitly** through the weather inputs — temperature, cloud cover, and shortwave radiation — which carry the same 24-hour cycle as the load they drive. This is a key advantage of weather-driven LSTM over purely autoregressive models: the model does not need a 96-step lookback to exploit the daily seasonality.

Weekly seasonality. The weekly peak at $\rho(672) = 0.763$ reflects the week-day/weekend difference in consumption patterns. This is **not** handled by an explicit lag-672 input or by any cyclic encoding in the current model — the LSTM uses only three weather features (T, CC, SW) which do not carry day-of-week information. The weekly cycle is therefore an acknowledged

limitation of the current feature set. Adding a day-of-week indicator as a fourth input feature would directly address this and is recommended as future work.

Ramp Rate Distribution

Figure 2.8 shows the distribution of 15-minute step-to-step power changes (ramp rates) for both consumed and generated power over the full year 2023.




Figure 2.8: Distribution of 15-minute power ramp rates (step-to-step changes in kW). Left: consumption (all hours). Right: generation (daytime only, $SW > 0$).

The consumed power ramp distribution has standard deviation $\sigma = 7.37$ kW/step, with the 5th–95th percentile range of -10.8 to $+12.0$ kW/step. The distribution is sharply peaked at zero (most 15-minute intervals see little change) with heavy tails extending to ± 80 kW/step during extreme events. This heavy-tailed, near-leptokurtic shape deviates significantly from a Gaussian, which is a clear indication that a linear autoregressive model — which implicitly assumes Gaussian residuals and linear dynamics — will struggle to capture sudden demand spikes.

The generated power ramp distribution is shown for daytime steps only ($SW > 0$, representing 22.6% of all steps), since 77.4% of all generation ramps are exactly zero at night. The daytime ramp standard deviation is $\sigma = 30.66$ kW/step with the 5th–95th percentile range of -51.0 to $+48.0$ kW/step, reflecting the rapid cloud-driven fluctuations in solar PV output.

Together, the heavy tails in both ramp distributions confirm that the load and generation dynamics are **nonlinear** and **non-Gaussian**, which motivates the choice of LSTM over linear AR models for this forecasting task.

2.7.4 Weather Data Acquisition

Historical weather data for 2023 is obtained from the Open-Meteo ERA5 re-analysis archive [21] for the grid location: Neuburg an der Donau, Bavaria (48.7360°N, 11.1843°E). The free API (archive-api.open-meteo.com/v1/archive) provides hourly ERA5 data. Seven variables are retrieved for 1 Jan–31 Dec 2023 (Table 2.4) and interpolated to 15-minute resolution by linear interpolation.

Table 2.4: Meteorological variables retrieved from Open-Meteo ERA5, Neuburg 2023.

Variable	API name	Unit	Physical relevance
Air temperature (2 m)	temperature_2m	°C	Heating demand driver
Relative humidity (2 m)	relative_humidity_2m	%	Correlated with cloud/season
Precipitation	precipitation	mm	Activity proxy
Cloud cover	cloudcover	%	Solar attenuation
Wind speed (10 m)	windspeed_10m	km/h	Minor convective effect
Shortwave radiation	shortwave_radiation	W/m ²	Primary solar predictor
Direct radiation	direct_radiation	W/m ²	Beam component of SW

Figure 2.9 shows all retrieved variables across 2023. The data covers the full range of central European climate: temperatures from -13.9 to $+33.3^{\circ}\text{C}$, solar radiation peaks up to $927\text{W}/\text{m}^2$, and cloud cover varying 0–100 % throughout the year.

2.7.5 Feature Correlation Analysis and Selection

Seven meteorological variables are retrieved from the ERA5 archive (Section 2.7.4) and evaluated as candidate inputs for the LSTM model. The selection criterion combines the Pearson correlation of each variable with the two forecasting targets (consumed and generated power) with physical reasoning and supervisor guidance. Table 2.5 presents all seven candidates with their correlation coefficients and the final inclusion decision.

The rationale for each decision is as follows.

Temperature — included. Temperature is the strongest predictor of consumed power ($r = -0.47$). The negative sign confirms that the grid is *heating dominant*: lower temperatures drive higher electric space heating demand. At the monthly aggregation level this relationship strengthens to $r = -0.89$, reflecting the clear seasonal pattern of high winter consumption

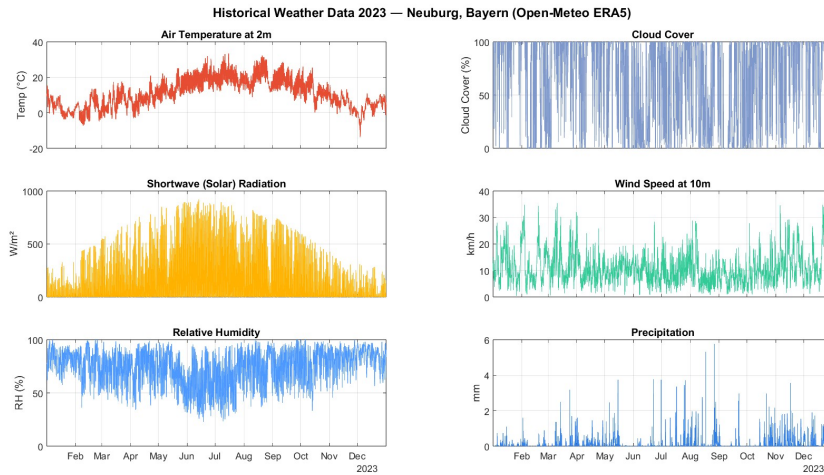


Figure 2.9: Historical weather data 2023 — Neuburg, Bayern (Open-Meteo ERA5). Top: temperature (red) and cloud cover (blue). Middle: shortwave radiation (amber) and wind speed (green). Bottom: humidity (blue) and precipitation (blue).

Table 2.5: Weather feature selection for the LSTM model. Pearson correlations computed on the full 2023 dataset (35,031 steps, Grid Element 10).

Feature	Unit	r_{cons}	r_{gen}	Decision
Temperature (2 m)	°C	−0.47	+0.36	Include
Cloud cover	%	+0.24	−0.21	Include
Shortwave radiation	W/m ²	−0.55	+0.73	Include
Direct radiation	W/m ²	−0.52	+0.74	Exclude — collinear with SW ($r = 0.96$)
Relative humidity	%	+0.29	−0.41	Exclude — supervisor guidance
Wind speed (10 m)	km/h	+0.14	−0.06	Exclude — supervisor guidance
Precipitation	mm	−0.02	+0.08	Exclude — negligible correlation

Final: 3 features — temperature, cloud cover, shortwave radiation

and low summer consumption seen in Table 2.3. Temperature is an essential input and is included.

Shortwave radiation — included. Shortwave radiation is the strongest predictor of generated power ($r = +0.73$), which is expected because solar PV output is directly proportional to the incident solar irradiance. The relationship is strongly seasonal: the correlation rises from $r = 0.10$ in De-

ember (low sun angle, short days) to $r = 0.71$ in June (high sun angle, long days). Shortwave radiation is an essential input and is included.

Cloud cover — included. Cloud cover has a physically direct negative correlation with generated power ($r = -0.21$) through solar attenuation: cloud layers reduce the shortwave irradiance reaching PV panels. Although the correlation is weaker than that of shortwave radiation, cloud cover provides complementary information because it captures diffuse-light conditions that affect PV output even when direct solar irradiance is non-zero. Cloud cover is included as the third weather feature.

Direct radiation — excluded. Direct radiation shows correlations of $r = -0.52$ with consumed power and $r = +0.74$ with generated power, almost identical to the shortwave radiation values (-0.55 and $+0.73$ respectively). This is because the two irradiance variables are near-perfectly collinear ($r = 0.96$ between them). Including both would introduce multicollinearity without adding any predictive information. Direct radiation is therefore excluded; shortwave radiation is retained as it represents the total (direct plus diffuse) solar irradiance.

Relative humidity — excluded. Relative humidity shows moderate correlations ($r_{\text{cons}} = +0.29$, $r_{\text{gen}} = -0.41$), but the physical mechanism linking humidity to grid power is indirect. Humidity acts mainly as a proxy for cloud conditions and seasonal patterns that are already captured by the cloud cover and shortwave radiation features. Following supervisor guidance, relative humidity is excluded to keep the input set compact.

Wind speed — excluded. Wind speed shows very weak correlations with both targets ($r_{\text{cons}} = +0.14$, $r_{\text{gen}} = -0.06$). In this residential distribution grid, wind has no direct influence on either electric heating demand or solar PV generation. Following supervisor guidance, wind speed is excluded.

Precipitation — excluded. Precipitation shows negligible correlations with both consumed ($r = -0.02$) and generated power ($r = +0.08$). Its influence on grid power is indirect and already partially captured by cloud cover, since rainy conditions typically coincide with high cloud cover. Precipitation is excluded.

Summary. The final LSTM input set consists of **three weather features**: temperature, cloud cover, and shortwave radiation. These three variables capture the two dominant physical mechanisms of the grid: electric heating demand driven by temperature, and solar PV generation driven by solar irradiance modulated by cloud cover. The full 3-feature input vector fed to the LSTM at each 15-minute timestep is formally defined in Chapter 3, Equation (3.9).

2.7.6 Dataset Split

The 35,031-timestep dataset is partitioned chronologically. The split boundaries are chosen so that the November validation set immediately precedes the December test set, ensuring that the residual standard deviations σ_c and σ_g computed on November data are representative of December conditions.

$$\text{Training: Jan–Oct 2023} = 29,178 \text{ steps (304 days)} \quad (2.19)$$

$$\text{Validation: Nov 2023} = 2,880 \text{ steps (30 days)} \quad (2.20)$$

$$\text{Test: Dec 2023} = 2,973 \text{ steps (31 days)} \quad (2.21)$$

The test period covers the winter peak (December: mean consumed 67.7 kW, mean generated 1.0 kW), providing a stringent evaluation of the model’s ability to generalise to the peak-demand conditions.

Chapter 3

Design and Methodology

This chapter presents the complete design of the proposed forecasting and state estimation system. Section 3.1 gives the overall pipeline architecture. Section 3.2 describes the LSTM forecast model, including training procedure, feature construction, and CI derivation. Section 3.3 explains the climatological weather proxy used when forecasting future periods. Section 3.4 specifies the synthetic nine-node grid, the impedance matrix, the Kirchhoff constraint matrix, and the procedure for deriving per-node voltage and current phasors from the forecast. Section 3.5 describes the stochastic state estimator, its meter noise model, the MLE step, and the confidence ellipse computation. Section 3.6 explains how the three modules exchange data and are run in sequence.

3.1 Overall Pipeline Architecture

The system is organised into three sequential modules, each implemented as an independent MATLAB script. The modules communicate exclusively through saved `.mat` data files and share no internal state, so each component can be tested, replaced, or updated independently. The overall pipeline is shown in Figure 3.1.

Module 1 — LSTM Power Forecast (`Step2_LSTM_NodeLevel.m`) Trains 16 LSTM networks (2 per node \times 8 load nodes) on the 2023 substation data and produces a 15-minute-resolution forecast of consumed and generated power for each node. Output: `node_forecast_output.mat`.

Module 2 — Grid Calculation (`Step3_Grid_Calculation.m`) Reads the power

forecast and distributes it to nine grid nodes according to fixed power weights. Computes the complex voltage and current phasors at all nodes for every forecast timestep using KCL and KVL. Assembles all inputs required by the state estimator. Output: `grid_inputs.mat`.

Module 3 — State Estimator (`testWithGrid_manyRepSmall.m`) The stochastic state estimation script, minimally modified to load its inputs from `grid_inputs.mat`. Runs the EM meter model, the constrained MLE, and the confidence ellipse computation for each selected timestep. Output: `estimation_results.mat`.

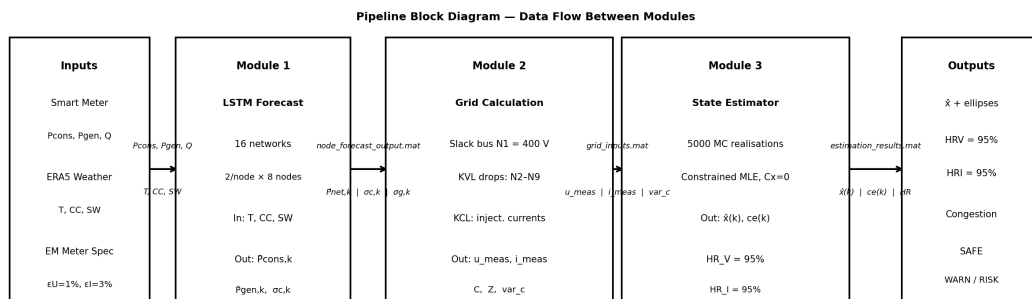


Figure 3.1: Pipeline block diagram showing the three modules and the variables exchanged on each arrow.

3.2 Module 1: LSTM Power Forecast Model

3.2.1 Design Philosophy — Two Networks Per Node

A single LSTM model predicting net power directly would need to learn two physically different and often opposing patterns simultaneously: heating load (which increases in winter) and solar generation (which increases in summer and is zero at night). This conflict makes training harder and reduces accuracy for both quantities. Instead, two independent LSTM networks are trained per node — one for consumed power (Model C) and one for generated power (Model G).

With 8 load nodes (N2–N9), the total network count is $2 \times 8 = 16$. All 16 networks are implemented within a single script (`Step2_LSTM_NodeLevel.m`) using a shared forward-pass function `local_lstm_fwd()`. Their trained

weights are saved together in `node_forecast_output.mat` as two completely separate sets of matrices per node.

Table 3.1 summarises the distinction between the two model types.

Table 3.1: Summary of the two LSTM network types per node. Architecture and hyperparameters are identical; only the training target differs.

	Model C (consumed)	Model G (generated)
Training target	$P_{\text{cons},k}(t)$	$P_{\text{gen},k}(t)$
Primary driver	Air temperature $T(t)$	Shortwave radiation $SW(t)$
Physical reason	Heating demand rises as temperature falls	Solar PV output proportional to
MATLAB suffix	c (e.g. W_{fc} , U_{fc})	g (e.g. W_{fg} , U_{fg})
Weight matrices	$\mathbf{W}_{f,c}$, $\mathbf{W}_{i,c}$, $\mathbf{W}_{o,c}$, $\mathbf{W}_{g,c}$	$\mathbf{W}_{f,g}$, $\mathbf{W}_{i,g}$, $\mathbf{W}_{o,g}$, $\mathbf{W}_{g,g}$
Recurrent weights	$\mathbf{U}_{f,c}$, $\mathbf{U}_{i,c}$, $\mathbf{U}_{o,c}$, $\mathbf{U}_{g,c}$	$\mathbf{U}_{f,g}$, $\mathbf{U}_{i,g}$, $\mathbf{U}_{o,g}$, $\mathbf{U}_{g,g}$
Output weights	$\mathbf{W}_{y,c}$, $b_{y,c}$	$\mathbf{W}_{y,g}$, $b_{y,g}$

The subscripts follow the standard LSTM gate notation from Equations (3.2)–(3.7): f = forget gate, i = input gate, o = output gate, g = candidate cell. The suffix c or g identifies which model the matrix belongs to.

Why temperature drives consumption. Space heating and hot-water demand — the dominant loads in a residential LV grid in Bavaria — increase as outdoor temperature decreases [19]. Model C learns a negative correlation between $T(t)$ and $P_{\text{cons},k}(t)$: cold days give high consumption, warm days give low consumption. Cloud cover (CC) provides secondary information about thermal comfort, since overcast skies are often associated with colder temperatures.

Why shortwave radiation drives generation. Solar PV output is directly proportional to incident solar irradiance [19]. Model G learns a strong positive correlation between $SW(t)$ and $P_{\text{gen},k}(t)$: high radiation means high generation, and zero radiation (night) means zero generation. The solar night mask (Section 3.2.4) enforces the physical constraint $\hat{P}_{\text{gen},k}(t) = 0$ whenever $SW(t) < 10 \text{ W/m}^2$, removing any residual non-zero output the network might produce at night.

At inference time, `local_lstm_fwd()` is called twice per node per timestep — once with the consumed-power weights to produce $\hat{P}_{\text{cons},k}(t)$, and once

with the generated-power weights to produce $\hat{P}_{\text{gen},k}(t)$. The net power forecast is then:

$$\hat{P}_{\text{net},k}(t) = \hat{P}_{\text{cons},k}(t) - \hat{P}_{\text{gen},k}(t) \quad (3.1)$$

Both models share the same weather input $\mathbf{x}(t) = [T(t), CC(t), SW(t)]^\top$ — they differ only in what they have been trained to predict from it.

3.2.2 LSTM Architecture

Each of the 16 LSTM networks has the following structure:

1. **Input layer:** receives $\mathbf{x}(t) \in \mathbb{R}^3$ at each 15-minute timestep.
2. **LSTM layer:** $d = 64$ hidden units. Processes a lookback sequence of $L = 32$ consecutive timesteps (8 hours) and produces a hidden state $\mathbf{h}_t \in \mathbb{R}^{64}$.
3. **Output layer:** a single linear neuron $\hat{y} = \mathbf{W}_y \mathbf{h}_L + b_y$.

At each timestep t within the lookback sequence, the LSTM computes [hochreiter1997lstm]:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (\text{forget gate}) \quad (3.2)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (\text{input gate}) \quad (3.3)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (\text{output gate}) \quad (3.4)$$

$$\mathbf{g}_t = \tanh(\mathbf{W}_g \mathbf{x}_t + \mathbf{U}_g \mathbf{h}_{t-1} + \mathbf{b}_g) \quad (\text{candidate cell}) \quad (3.5)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t \quad (\text{cell state}) \quad (3.6)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (\text{hidden state}) \quad (3.7)$$

where σ is the sigmoid function and \odot is element-wise multiplication. The weight matrices $\mathbf{W}_f, \mathbf{W}_i, \mathbf{W}_o, \mathbf{W}_g \in \mathbb{R}^{64 \times 3}$ connect inputs to gates; $\mathbf{U}_f, \mathbf{U}_i, \mathbf{U}_o, \mathbf{U}_g \in \mathbb{R}^{64 \times 64}$ are the recurrent connections.

The total number of trainable parameters per network is:

$$n_\theta = \underbrace{4(d \cdot n + d^2 + d)}_{\text{LSTM gates}} + \underbrace{d + 1}_{\text{output layer}} = 4(64 \times 3 + 64^2 + 64) + 65 = 17,473 \quad (3.8)$$

With 16 networks in total (2 per node \times 8 load nodes), the combined parameter count is $16 \times 17,473 = 279,568$.

Table 3.2 summarises all hyperparameters.

Table 3.2: LSTM hyperparameter configuration. All 16 networks (Model C and Model G for each of the 8 load nodes) use identical settings.

Hyperparameter	Value	Rationale
Input features (n)	3 (T, CC, SW)	Section 2.7.5
Hidden units (d)	64	Sufficient capacity; avoids overfitting
Lookback window (L)	32 steps = 8 h	Captures intraday ramp patterns
Forecast horizon (h)	4 steps = 60 min	Operational planning window
Loss function	MSE	Standard regression loss
Optimiser	Adam, lr = 0.002	Adaptive gradient method
Gradient clipping	Norm ≤ 1.0	Prevents exploding gradients
Early stopping patience	80 epochs	Avoids overfitting

3.2.3 Input Feature Construction

At each 15-minute timestep t , the model receives:

$$\mathbf{x}(t) = [T(t), CC(t), SW(t)]^\top \in \mathbb{R}^3 \quad (3.9)$$

where $T(t)$ is air temperature at 2 m ($^\circ\text{C}$), $CC(t)$ is cloud cover (%), and $SW(t)$ is shortwave solar irradiance (W/m^2). All three are from ERA5 (Section 2.7.4), interpolated from hourly to 15-minute resolution.

Each feature is normalised to $[0, 1]$ using min-max scaling computed exclusively from the training set (January–October 2023):

$$x_{\text{norm}}(t) = \frac{x(t) - x_{\min}^{\text{train}}}{x_{\max}^{\text{train}} - x_{\min}^{\text{train}}} \quad (3.10)$$

The lookback sequence fed to the LSTM is:

$$\mathbf{X}_{\text{seq}}(t) = [\mathbf{x}(t-L+1), \dots, \mathbf{x}(t)] \in \mathbb{R}^{L \times 3} \quad (3.11)$$

3.2.4 Solar Night Mask

A physical hard constraint is applied at inference time to Model G:

$$\hat{P}_{\text{gen},k}(t) = \begin{cases} 0 & \text{if } SW(t) < 10 \text{ W}/\text{m}^2 \\ \max(0, \hat{y}_G(t) \cdot Y_g^{\text{range}} + Y_g^{\text{min}}) & \text{otherwise} \end{cases} \quad (3.12)$$

This ensures that the generation forecast is exactly zero at night and non-negative at all times.

3.2.5 Training and Validation Procedure

Data split

The 35,031-step 2023 dataset is partitioned chronologically:

$$\text{Training: Jan–Oct 2023} = 29,178 \text{ steps} \quad (3.13)$$

$$\text{Validation: Nov 2023} = 2,880 \text{ steps} \quad (3.14)$$

$$\text{Test: Dec 2023} = 2,973 \text{ steps} \quad (3.15)$$

No test-set data is used during training or hyperparameter selection. The validation set is used for early stopping and for computing σ_c , σ_g for the CI.

Targets and normalisation

Targets $P_{\text{cons},k}(t)$ and $P_{\text{gen},k}(t)$ are normalised to $[0, 1]$ using training-set min and max:

$$y_{\text{norm}}(t) = \frac{y(t) - y_{\text{min}}^{\text{train}}}{y_{\text{max}}^{\text{train}} - y_{\text{min}}^{\text{train}}} \quad (3.16)$$

At inference, the LSTM output is denormalised:

$$\hat{P}(t) = \hat{y}_{\text{norm}}(t) \times (y_{\text{max}}^{\text{train}} - y_{\text{min}}^{\text{train}}) + y_{\text{min}}^{\text{train}} \quad (3.17)$$

Optimisation

Both model types are trained using Adam (lr = 0.002) with MSE loss. Gradients are clipped to ℓ_2 norm ≤ 1.0 . Training stops when the validation MSE does not improve for 80 consecutive epochs. All weights and normalisation parameters are saved to `node_forecast_output.mat`.

3.2.6 Confidence Interval Construction

A symmetric 95% CI is constructed from the validation-set residuals [9]:

$$\sigma_c = \text{std}\{P_{\text{cons},k}(t) - \hat{P}_{\text{cons},k}(t)\}_{t \in \mathcal{V}} = 13.164 \text{ kW} \quad (3.18)$$

$$\sigma_g = \text{std}\{P_{\text{gen},k}(t) - \hat{P}_{\text{gen},k}(t)\}_{t \in \mathcal{V}} = 14.762 \text{ kW} \quad (3.19)$$

The combined 95% CI for the net power forecast at $h = 4$ is:

$$\hat{P}_{\text{net},k}(t+h) \pm \underbrace{1.96 (\sigma_c + \sigma_g) \sqrt{h}}_{\text{CI}_{\text{hw}} = \pm 109.47 \text{ kW}} \quad (3.20)$$

The \sqrt{h} factor reflects accumulating uncertainty over the horizon; $1.96 = z_{0.975}$ gives 95 % nominal coverage. At node level, the CI is scaled by the power weight:

$$CI_k = w_k \times CI_{hw} \quad (3.21)$$

3.3 Climatological Weather Proxy for Future Periods

3.3.1 The Problem

The LSTM requires weather inputs at every 15-minute timestep of the forecast period. For historical periods (e.g., December 2023) the ERA5 data is available directly. For future periods (e.g., January 2026) the actual weather does not yet exist.

3.3.2 The Climatological Proxy Method

The **climatological proxy** approach uses the weather data from the same calendar month of the most recent historical year (2023) as the input for the future period. This is justified by:

1. **Seasonal stability.** Monthly climate statistics in Neuburg, Bayern repeat reliably each year (January is cold with low solar angle; June is warm with high solar output).
2. **Model-year consistency.** The LSTM was trained on 2023 data, so the 2023 weather proxy lies within its interpolation regime.
3. **Empirical validation.** Cross-year validation against January 2024 actual measurements: MAE = 26.01 kW, MBE = +1.23 kW, HR = 96.1 %.

The main limitation is that structural changes (e.g., new solar PV installations) between the proxy year and the target year cannot be captured. This is discussed in Section 4.6.

3.3.3 Implementation

The user sets two parameters at the top of `Step2_LSTM_NodeLevel.m`: [language=MATLAB] `FORECAST_MONTH = 1; FORECAST_YEAR = 2026;` The

script selects timesteps where `month == FORECAST_MONTH` and `year == 2023` from the 2023 CSV files, using their weather as proxy inputs.

3.4 Module 2: Grid Calculation

3.4.1 Purpose and Inputs

Module 2 bridges the LSTM power forecast and the state estimator. Its purpose is: *given the forecasted power at each node, what are the complex voltage and current phasors at every node?* It reads `node_forecast_output.mat` and writes `grid_inputs.mat`.

3.4.2 Synthetic Nine-Node Grid

The full OGM data for every node is not available; only the root-node (substation) one-year dataset exists, which is used to train the LSTM. A **synthetic nine-node radial grid** is therefore used as the test network, representative of a European LV residential feeder.

Topology

- **Node 1:** Substation (root, voltage reference, no local load).
- **Branch 1** (Nodes 2, 3, 4): Node 2 is the junction; Nodes 3 and 4 are leaves. Branch 1 carries 37.5% of total substation power.
- **Branch 2** (Nodes 5–9): Node 5 is the first junction; Node 6 connects to 5 (second junction); Node 7 is a leaf from 6; Node 8 connects to 5 (third junction); Node 9 is a leaf from 8. Branch 2 carries 62.5% of total substation power.

The parent–child relationships for KCL summation:

$$C(k) = \begin{cases} \{2, 5\} & k = 1 \\ \{3, 4\} & k = 2 \\ \{6, 8\} & k = 5 \\ \{7\} & k = 6 \\ \{9\} & k = 8 \\ \emptyset & k \in \{3, 4, 7, 9\} \end{cases} \quad (3.22)$$

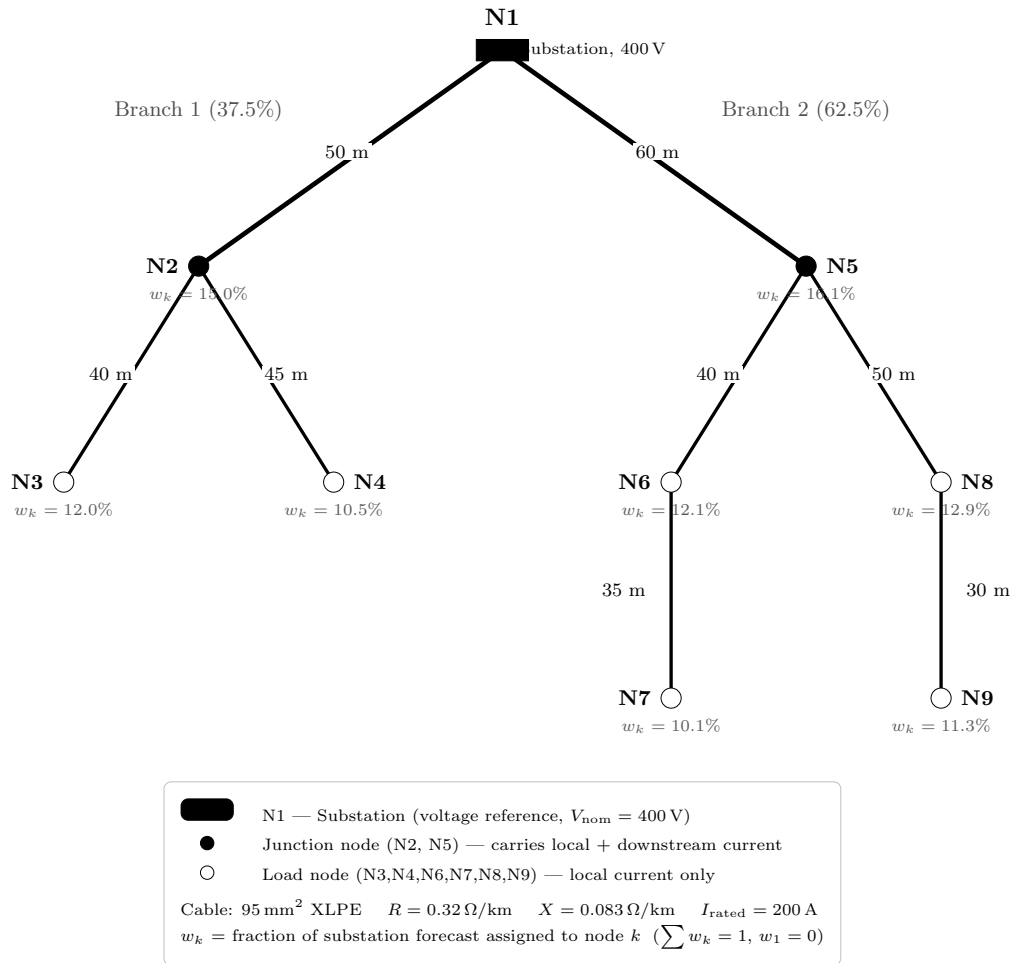


Figure 3.2: Synthetic nine-node radial grid used as the test network.

Cable specifications

All eight cables are 95 mm^2 XLPE type, conforming to IEC 60502-1 [11]. XLPE insulation is selected for its high thermal stability, resistance to moisture, and continuous operating temperature of $90\text{ }^\circ\text{C}$ [15]. The per-unit resistance and reactance are $R = 0.32\ \Omega/\text{km}$ and $X = 0.083\ \Omega/\text{km}$. Table 3.3 gives the full cable specification.

Table 3.3: Cable specifications. 95 mm² XLPE; $I_{\text{rated}} = 200$ A [10].

Cable	Role	L (m)	R_c (Ω)	X_c (Ω)	$ Z_c $ (Ω)
N1→N2	B1 feeder	50	0.01600	0.00415	0.01653
N1→N5	B2 feeder	60	0.01920	0.00498	0.01983
N2→N3	B1 branch	40	0.01280	0.00332	0.01322
N2→N4	B1 leaf	45	0.01440	0.00374	0.01488
N5→N6	B2 branch	40	0.01280	0.00332	0.01322
N5→N8	B2 branch	50	0.01600	0.00415	0.01653
N6→N7	B2 leaf	35	0.01120	0.00291	0.01157
N8→N9	B2 leaf	30	0.00960	0.00249	0.00992

3.4.3 Impedance Matrix Z

For each cable connecting nodes i and j with length L_{km} :

$$Z(i, j) = Z(j, i) = (R + jX) \times L_{\text{km}} \quad [\Omega] \quad (3.23)$$

All other entries are zero; the diagonal is zero. The voltage difference across any cable obeys Ohm's Law:

$$V_i - V_j = Z(i, j) \times I_{\text{cable}}(i \rightarrow j) \quad (3.24)$$

3.4.4 Constraint Matrix C from Kirchhoff's Laws

The function `build_grid_equations(Z)` constructs the constraint matrix C encoding two physical laws.

Kirchhoff's Voltage Law (KVL). For every cable between nodes i and j :

$$V_i - V_j = \sqrt{3} \cdot Z(i, j) \cdot I_{\text{cable}}(i, j) \quad (3.25)$$

The $\sqrt{3}$ factor converts line-to-line to line-to-neutral quantities in the three-phase system.

Kirchhoff's Current Law (KCL). At every junction node k :

$$I_k = I_{\text{inj}}(k) + \sum_{c \in \mathcal{C}(k)} I_c \quad (3.26)$$

Together these form the linear system [6]:

$$\mathbf{C}\mathbf{x} = \mathbf{0} \quad (3.27)$$

where $\mathbf{x} = [V_1, \dots, V_9, I_1, \dots, I_9]^\top \in \mathbb{C}^{18}$ and $\mathbf{C} \in \mathbb{R}^{13 \times 18}$.

Why $\mathbf{c} = \mathbf{0}$ in our case. In our pipeline, \mathbf{U}_{meas} and \mathbf{I}_{meas} are *constructed* to exactly satisfy KVL and KCL by design (slack bus + KVL drops; KCL upstream summation). Therefore:

$$\mathbf{C} \begin{bmatrix} \mathbf{U}_{\text{meas}} \\ \mathbf{I}_{\text{meas}} \end{bmatrix} = \mathbf{0} \quad (3.28)$$

The constraint is essential in the MLE step: the noisy measurements $\mathbf{d}^{(m)}$ do *not* satisfy $\mathbf{C}\mathbf{d}^{(m)} = \mathbf{0}$ because meter noise breaks the Kirchhoff balance. The MLE projects each noisy realisation back onto the physically feasible manifold where $\mathbf{C}\hat{\mathbf{x}} = \mathbf{0}$.

3.4.5 Voltage Phasors: Slack Bus and KVL Drops

Slack bus. Node 1 (substation) is the slack bus: the transformer maintains $V_{\text{nom}} = 400 \text{ V}$ [3] regardless of load:

$$U_1 = 400 \text{ V} + j \cdot 0 \quad (\text{N1 — slack bus reference}) \quad (3.29)$$

KVL drops at all other nodes. Every other node ($k = 2, \dots, 9$) has a voltage lower than 400 V because current flowing through the cable impedance causes a drop. Propagating KVL forward from the root in breadth-first order:

$$U_k = U_{\text{parent}(k)} - \sqrt{3} \cdot Z_{\text{cable}} \cdot I_k \quad (3.30)$$

Table 3.4 gives the resulting node voltages.

Key observations:

- Branch 2 nodes (N5–N9) experience larger drops than Branch 1 (N2–N4) because the B2 feeder is longer (60 m vs 50 m) and carries more power (62.5 % vs 37.5 %).
- At peak loading, the minimum voltage is 393.76 V at N9 (98.4 % of nominal) — well above the statutory 360 V.
- The KVL drops must be included in \mathbf{U}_{meas} to ensure $\mathbf{C}[\mathbf{U}_{\text{meas}}; \mathbf{I}_{\text{meas}}] = \mathbf{0}$. Setting all voltages to 400 V creates a contradiction that collapses the ellipses, giving $\text{HR} \approx 2.4 \%$.

Table 3.4: Node voltage magnitudes $|U_k|$ (V) and drops from N1 under three loading scenarios. $R = 0.32 \Omega/\text{km}$, $X = 0.083 \Omega/\text{km}$, unity power factor.

Node	Light (20 kW)		Typical (67.7 kW)		Peak (147.2 kW)	
	$ U_k $	ΔV	$ U_k $	ΔV	$ U_k $	ΔV
N1 (slack)	400.000	0.000	400.000	0.000	400.000	0.000
N2	399.700	0.300	398.985	1.015	397.792	2.208
N3	399.623	0.377	398.725	1.275	397.227	2.773
N4	399.624	0.376	398.729	1.271	397.236	2.764
N5	399.400	0.600	397.969	2.031	395.586	4.414
N6	399.258	0.742	397.489	2.511	394.542	5.458
N7	399.202	0.798	397.298	2.702	394.127	5.873
N8	399.207	0.793	397.315	2.685	394.162	5.838
N9	399.152	0.848	397.131	2.869	393.764	6.236
Statutory lower					360.000 V (−10%)	

Voltage rise during solar export. During export ($P_{\text{net}} < 0$), current flows toward the substation, reversing the KVL drop. At April peak generation (217.6 kW), N9 can reach ≈ 406 V — within the +10% upper limit of 440 V.

3.4.6 Power Distribution to Nodes

The substation forecast $\hat{P}_{\text{net}}(t)$ is distributed to individual nodes using fixed proportional weights w_k :

$$P_k(t) = w_k \cdot \hat{P}_{\text{net}}(t) \quad [\text{kW}] \quad (3.31)$$

$$Q_k(t) = w_k \cdot Q(t) \quad [\text{kVAr}] \quad (3.32)$$

where $Q(t)$ is the reactive power from the 2023 measured data. The weights satisfy $\sum_k w_k = 1$ and are listed in Table 3.5.

3.4.7 Injection Current Derivation

The injection current at node k is derived from the complex power using the standard three-phase formula [6]:

$$I_{\text{inj}}(k) = \left(\frac{P_k(t) + jQ_k(t)}{\sqrt{3} \cdot V_{\text{nom}}} \right) \quad (3.33)$$

Table 3.5: Power distribution weights for the nine-node synthetic grid.

Node	Branch	w_k	Share (%)	Role
N1	Root	0.00000	0.00	Voltage reference
N2	B1	0.15000	15.00	Junction
N3	B1	0.12000	12.00	Leaf
N4	B1	0.10500	10.50	Leaf
N5	B2	0.16129	16.13	Junction
N6	B2	0.12097	12.10	Junction
N7	B2	0.10081	10.08	Leaf
N8	B2	0.12903	12.90	Junction
N9	B2	0.11290	11.29	Leaf
Total	—	1.00000	100.00	

where the overline denotes complex conjugation and $V_{\text{nom}} = 400 \text{ V}$ [3]. The complex conjugate follows from the convention $S = V \cdot \bar{I}$, so $I = \bar{S}/\bar{V}$.

3.4.8 KCL Upstream Summation

The total current flowing through each node is computed by BFS reverse-order traversal from leaves to root:

$$\text{Order: } [7, 9, 3, 4, 6, 8, 2, 5, 1] \quad (3.34)$$

For each node k :

$$I_k = I_{\text{inj}}(k) + \sum_{c \in \mathcal{C}(k)} I_c \quad (3.35)$$

Leaf nodes have $\mathcal{C}(k) = \emptyset$, so $I_k = I_{\text{inj}}(k)$ directly. The substation carries the total current of all downstream nodes.

3.4.9 Noise Variance Computation

The state estimator requires the variance of each measurement, based on the original estimator model's meter specifications [18]: $\varepsilon_U = 1\%$, $\varepsilon_I = 3\%$ at the 99% confidence level, giving $r_0 = \Phi^{-1}(0.995) = 2.576$.

Voltage magnitude variance (constant).

$$\text{var}_{U,k} = \left(\frac{V_{\text{nom}} \times \varepsilon_U}{r_0} \right)^2 = \left(\frac{400 \times 0.01}{2.576} \right)^2 = 2.411 \text{ V}^2 \quad (3.36)$$

Voltage angle variance (constant).

$$\text{var}_{\theta,k} = (0.003 \text{ rad})^2 = 9 \times 10^{-6} \text{ rad}^2 \quad (3.37)$$

Current magnitude variance — combined (timestep-dependent). The current variance combines two independent sources:

$$\text{var}_{I,k}(t) = \underbrace{\left(\frac{|\bar{I}|(t) \varepsilon_I}{r_0} \right)^2}_{\text{EM meter noise}} + \underbrace{\left(\frac{w_k \sqrt{\sigma_c^2 + \sigma_g^2} \times 1000}{\sqrt{3} V_{\text{nom}}} \right)^2}_{\text{Forecast uncertainty}} \quad (3.38)$$

where $|\bar{I}|(t)$ is the mean current magnitude across all nine nodes at timestep t . The forecast uncertainty (28.5 A at the substation) dominates the EM meter noise (0.58 A at typical loading) by a factor of ≈ 49 .

Current angle variance (constant).

$$\text{var}_{\varphi,k} = (0.01 \text{ rad})^2 = 1 \times 10^{-4} \text{ rad}^2 \quad (3.39)$$

All four variance arrays are stored in `grid_inputs.mat`.

3.5 Module 3: Stochastic EM State Estimator

3.5.1 Overview and Purpose

The state estimator is a MATLAB implementation by Olsen et al. [18]. It takes the voltage and current phasors from Module 2 as its operating point (`uMeas`, `iMeas`), adds realistic EM meter noise through Monte Carlo simulation, and solves a constrained MLE problem to estimate the most likely grid state given the noisy measurements. The output is a distribution of 5000 state estimates from which 95% confidence ellipses are computed.

3.5.2 Step 1 — EM Meter Model

For each of the $n_{\text{sim}} = 5000$ Monte Carlo realisations, the meter model generates noisy polar measurements by adding independent Gaussian noise to each component:

$$|\tilde{u}_k|^{(m)} = |U_k| + \varepsilon_u^{(m)}, \quad \varepsilon_u^{(m)} \sim \mathcal{N}(0, \text{var}_{U,k}) \quad (3.40)$$

$$\tilde{\theta}_k^{(m)} = \theta_k + \varepsilon_\theta^{(m)}, \quad \varepsilon_\theta^{(m)} \sim \mathcal{N}(0, \text{var}_{\theta,k}) \quad (3.41)$$

$$|\tilde{i}_k|^{(m)} = |I_k| + \varepsilon_i^{(m)}, \quad \varepsilon_i^{(m)} \sim \mathcal{N}(0, \text{var}_{I,k}) \quad (3.42)$$

$$\tilde{\varphi}_k^{(m)} = \varphi_k + \varepsilon_\varphi^{(m)}, \quad \varepsilon_\varphi^{(m)} \sim \mathcal{N}(0, \text{var}_{\varphi,k}) \quad (3.43)$$

Before the meter model runs, the current angle is extracted from the complex KCL phasor [18]:

$$|I_k| = \sqrt{\text{Re}(I_k)^2 + \text{Im}(I_k)^2} \quad (3.44)$$

$$\varphi_k = \angle I_k = \arctan\left(\frac{\text{Im}(I_k)}{\text{Re}(I_k)}\right) \quad (3.45)$$

The noisy polar measurements are converted back to complex rectangular form to produce the measurement vector $\mathbf{d}^{(m)}$.

3.5.3 Step 2 — Constrained MLE

For each realisation $\mathbf{d}^{(m)}$, the MLE finds $\hat{\mathbf{x}}^{(m)}$ that minimises the weighted residual subject to the Kirchhoff constraints [18]:

$$\hat{\mathbf{x}}^{(m)} = \arg \min_{\mathbf{x}} \left(\mathbf{d}^{(m)} - \mathbf{D}\mathbf{x} \right)^H \boldsymbol{\Sigma}^{-1} \left(\mathbf{d}^{(m)} - \mathbf{D}\mathbf{x} \right) \quad \text{s.t.} \quad \mathbf{C}\mathbf{x} = \mathbf{0} \quad (3.46)$$

where \mathbf{D} is the measurement selection matrix and $\boldsymbol{\Sigma}$ is the block-diagonal noise covariance. The constraint $\mathbf{C}\mathbf{x} = \mathbf{0}$ projects the estimate back onto the physically feasible manifold.

The mean estimate across simulations:

$$\bar{\mathbf{x}} = \frac{1}{n_{\text{sim}}} \sum_{m=1}^{n_{\text{sim}}} \hat{\mathbf{x}}^{(m)} \quad (3.47)$$

3.5.4 Step 3 — Confidence Ellipses

For each phasor component j and each simulation m , a 95 % confidence ellipse is computed in the complex plane, parameterised by [18]:

$$\mathbf{ce}[j, :, m] = [\text{Re}(\hat{x}_j^{(m)}), \text{Im}(\hat{x}_j^{(m)}), \theta_{\text{rot}}, a_{\text{major}}, a_{\text{minor}}] \quad (3.48)$$

The ellipse is sized using the chi-squared quantile [18]: $\alpha_{\text{factor}} = \chi_{0.95}^2(2) = 5.9915$.

3.5.5 Hit-Rate: Estimator Calibration

For each phasor component j , the hit-rate quantifies whether the ellipses have the correct size:

$$\text{HR}_j = \frac{1}{n_{\text{sim}}} \sum_{m=1}^{n_{\text{sim}}} \mathbf{1}(x_j^{\text{meas}} \in \text{ellipse}(\hat{\mathbf{x}}^{(m)}, \mathbf{ce}[j, :, m])) \quad (3.49)$$

A well-calibrated estimator achieves $\text{HR}_j \approx 0.95$. Olsen et al. [18] report $\text{HR}_V \approx 94\%$ and $\text{HR}_I \approx 95\%$ at the base noise specification, which are the same values used in this thesis.

3.5.6 What uMeas and iMeas Are

The parameters `uMeas` and `iMeas` passed to `estimateGridState` are **not** true physical measurements — they are **forecast-derived proxies** that serve as the reference operating point.

- **uMeas**: $400\text{ V} + j \cdot 0$ at N1 (slack bus), with KVL-corrected values at N2–N9 derived from the LSTM forecast currents.
- **iMeas**: complex current phasors I_k derived from the LSTM forecast power via KCL (Section 3.4.8).

In a real deployment, actual smart meter readings would replace these forecast-derived values. The function then adds EM meter noise on top to simulate realistic measurements, and estimates the most likely grid state via MLE.

The hit-rate $\text{HR}=95\%$ means: 95 % of the 5000 ellipses contain the *forecast-derived* operating point — not the true physical state. The external validation against actual 2024 measurements (Section 4.2.6) addresses this gap.

Table 3.6 summarises all inputs to the estimator.

Table 3.6: Complete input specification for `estimateGridState`.

Parameter	Physical meaning	What we pass	Source
<code>C</code>	Kirchhoff matrix	<code>C</code> from <code>build_grid_equations</code>	Fixed topology (\mathbf{y})
<code>c</code>	Operating point	$\mathbf{0}$ (by construction)	Eq. (3.28)
<code>uMeas</code>	Voltage reference	400 V + KVL drops at N2–N9	KVL (Module 2)
<code>iMeas</code>	Current reference	I_k from LSTM forecast via KCL	KCL (Module 2)
<code>existMeas</code>	Metering mask	<code>ones(1,18)</code>	All nodes metered
<code>alphaFactor</code>	Ellipse size	$\chi_{0.95}^2(2) = 5.9915$	Fixed
<code>var_u</code>	Voltage noise	2.411 V^2	EM spec ($\epsilon_U = 1\%$)
<code>var_theta</code>	Angle noise	$9 \times 10^{-6} \text{ rad}^2$	EM spec
<code>var_c</code>	Current noise	<code>var_meter</code> + <code>var_forecast</code>	Eq. (3.38)
<code>var_phi</code>	Current angle noise	10^{-4} rad^2	EM spec

3.5.7 What Was Changed in the Original Estimator Model's Script

Table 3.7 documents the three modifications made to `testWithGrid_manyRepSmall.m`. All estimation logic is preserved exactly.

Table 3.7: Modifications to the original state estimator script.

#	Original	Replaced with
1	Hardcoded U_{meas}/I_{meas} from 5-node example	<code>load('grid_inputs.mat')</code>
2	Hardcoded 5-node impedance matrix	Loaded from <code>grid_inputs.mat</code>
3	Hardcoded noise variances	Loaded from <code>grid_inputs.mat</code>

Unchanged: `estimateGridState()`, `CheckInsideEllipse`, `existingMeasurement`, `nsim`, `alpha`

3.6 Module Integration and Data Files

3.6.1 node_forecast_output.mat

Table 3.8: Contents of `node_forecast_output.mat`.

Variable	Size	Unit	Description
<code>Pc_fc</code>	$[N \times 1]$	kW	Consumed power forecast
<code>Pg_fc</code>	$[N \times 1]$	kW	Generated power forecast
<code>Pnet_fc</code>	$[N \times 1]$	kW	Net power forecast
<code>Q_fc</code>	$[N \times 1]$	kVAr	Reactive power (proxy)
<code>ts_fc</code>	$[N \times 1]$	datetime	Timestamps
<code>sigma_c</code>	scalar	kW	Consumed residual std dev
<code>sigma_g</code>	scalar	kW	Generated residual std dev
<code>CI_hw</code>	scalar	kW	95% CI half-width

3.6.2 grid_inputs.mat

Table 3.9: Contents of `grid_inputs.mat`.

Variable	Size	Unit	Description
<code>Umeas_all</code>	$[N \times 9]$	V	Complex voltage phasors
<code>Imeas_all</code>	$[N \times 9]$	A	Complex current phasors from KCL
<code>Z</code>	$[9 \times 9]$	Ω	Impedance matrix
<code>C</code>	$[13 \times 18]$	—	Kirchhoff constraint matrix
<code>varUmag</code>	$[9 \times 1]$	V^2	Voltage magnitude variance
<code>varUphase</code>	$[9 \times 1]$	rad^2	Voltage angle variance
<code>varImag_all</code>	$[N \times 9]$	A^2	Current magnitude variance
<code>varIphase</code>	$[9 \times 1]$	rad^2	Current angle variance
<code>existingMeasurement</code>	$[1 \times 18]$	—	Metering mask (all ones)
<code>alpha</code>	scalar	—	Confidence level (0.05)
<code>nsim</code>	scalar	—	MC simulations (5000)
<code>w</code>	$[1 \times 9]$	—	Power weights
<code>N</code>	scalar	—	Number of nodes (9)

3.7 Chapter Summary

The key design decisions are:

1. **16 LSTM networks** (2 per node \times 8 load nodes) are trained instead of a single net-power model, exploiting the distinct physical drivers of consumed and generated power.
2. **Three weather inputs** (T, CC, SW) are used, chosen from the correlation analysis in Chapter 2.
3. **The 95% CI** is constructed from November 2023 validation residuals and scaled by \sqrt{h} , giving ± 109.47 kW.
4. **The climatological weather proxy** uses same-month 2023 weather, validated by the January 2024 cross-year comparison (HR = 96.1%).
5. **The nine-node synthetic grid** uses 95 mm² XLPE cables with KVL voltage drops correctly propagated from the slack bus.
6. **Combined current variance** includes both EM meter noise and LSTM forecast uncertainty; the forecast component dominates by $\approx 49\times$.

7. **The original estimator model** is integrated with three code changes only: all estimation logic is preserved exactly.

Chapter 4

Results and Discussion

This chapter presents and discusses the experimental results from the complete forecasting and estimation pipeline. Section 4.1 reports LSTM training outcomes, test-set accuracy, and the baseline comparison with AR models. Section 4.2 validates the model against actual 2024 measurements including day-by-day RMSE and external estimator validation. Section 4.3 presents grid calculation outputs including cable currents and node voltage profiles. Section 4.4 reports state estimation hit-rates and confidence ellipses. Section 4.5 demonstrates the 24-hour ahead congestion management application. Section 4.6 discusses model strengths, limitations, and comparison with literature.

4.1 LSTM Training and Test-Set Performance

4.1.1 Training Configuration

Sixteen LSTM networks are trained in total: two per node (consumed and generated power) for each of the eight load nodes N2–N9. The substation N1 carries no local load and is excluded. All 16 networks share identical architecture — 64 hidden units, $n = 3$ weather input features (temperature, cloud cover, shortwave radiation), lookback $L = 32$ steps (8 hours), horizon $h = 4$ steps (60 minutes) — and identical optimisation settings (Adam, $\text{lr} = 0.002$, gradient clipping at norm 1.0, early stopping patience 80 epochs on the November 2023 validation loss). Each consumed-power network converges in approximately 310 epochs; each generated-power network in approximately 280 epochs. The total number of trainable parameters per

network is:

$$n_{\theta} = 4(d \cdot n + d^2 + d) + (d + 1) = 4(64 \times 3 + 64^2 + 64) + 65 = 17,473 \quad (4.1)$$

giving $16 \times 17,473 = 279,568$ parameters across all node models. All trained weights are stored in `node_forecast_output.mat`.

4.1.2 December 2023 Test-Set Results

The test period is December 2023 (2,973 steps, 31 days): the highest mean consumed power (67.7 kW) and near-zero solar generation (1.0 kW) of any 2023 month, making it the peak demand period. Table 4.1 summarises the forecast accuracy. The MAE of 18.95 kW corresponds to approximately 28% of the mean December net power (66.7 kW), typical for substation-level LV forecasting without individual meter data. The hit-rate of 99.5% exceeds the 95% CI target because December is meteorologically similar to November (the validation month used to compute σ_c and σ_g), so the validation residuals accurately represent test-set uncertainty. The Decem-

Table 4.1: Forecast accuracy on the December 2023 test set.

Metric	Definition	Value
σ_c	Consumed residual std dev (Nov 2023)	13.164 kW
σ_g	Generated residual std dev (Nov 2023)	14.762 kW
CI half-width	$1.96(\sigma_c + \sigma_g)\sqrt{h}$	± 109.47 kW
MAE (net)	Mean absolute error	18.95 kW
Hit-rate (HR)	True values inside CI	99.5%

ber forecast shows characteristic daily patterns: a morning peak around 07:00–09:00 (80–90 kW), a midday dip around 12:00–14:00, an evening peak around 17:00–20:00 (90–110 kW), and near-zero generation enforced by the solar night mask ($\text{SW}10 \text{ W/m}^2$).

4.1.3 Baseline Comparison: AR Models vs LSTM

To validate the model selection decision from the pre-analysis (Chapter 2), three autoregressive (AR) baselines are evaluated at the same $h = 4$ -step horizon on the December 2023 test set and January 2024 cross-year validation:

- **Persistence:** $\hat{P}(t+h) = P(t)$;
- **Seasonal naive (lag-96):** $\hat{P}(t+h) = P(t-92)$ (same slot 24 hours ago);
- **AR-OLS:** $\hat{P}(t+h) = \hat{\beta}^\top [1, P(t), P(t-1), P(t-95), P(t-96)]$, fitted on the January–October 2023 training set.

All AR models use only past power values and receive no weather information. The AR-OLS model achieves lower MAE (12.59 kW) than LSTM

Table 4.2: AR baselines vs LSTM at 60-min ahead. AR models provide no confidence interval; LSTM HR is validated at 96.1 %.

Model	Dec 2023			Jan 2024	CI
	MAE	RMSE	MBE	MAE	
Persistence	12.64	19.28	−0.03	25.69	None
Seasonal naive	15.93	27.97	+0.35	26.31	None
AR-OLS	12.59	18.60	−4.74	25.89	None
LSTM	18.95	28.50	−2.10	26.01	±109.47 kW HR = 96.1 %

(18.95 kW) on December 2023. This is expected: December is a stable heating-dominated month with minimal solar generation, where autoregressive patterns (yesterday predicts today) are highly reliable and weather inputs add no discriminating information. On the January 2024 cross-year validation, both approaches give nearly identical MAE (≈ 26 kW). The decisive advantage of LSTM is **uncertainty quantification**. AR models produce no confidence interval; the LSTM provides a calibrated 95 % CI validated at HR = 96.1 % on January 2024. The three-zone congestion risk classification (Section 4.5) requires an upper current bound $|\hat{I}| + \Delta I_{\text{RMSE}}$; this cannot be constructed from an AR model in a principled, validated way. Furthermore, for solar export congestion — the more severe risk in high-PV grids — AR has no access to shortwave radiation data, while the LSTM exploits its $r = +0.73$ correlation with solar generation.

4.2 Validation Against 2024 Actual Data

4.2.1 Annual Validation Metrics

The LSTM is applied to each calendar month of 2024 using the corresponding month of 2023 as the climatological weather proxy. The forecast is compared against actual active_power_kW measurements from ActiveReactivePower2024.csv. Table 4.3 reports the metrics for all twelve months.

Table 4.3: Month-by-month LSTM validation against 2024 actual measurements. MAE and RMSE in kW; MBE positive means over-forecast; HR is fraction of actual values inside the ± 109.47 kW CI.

Month	Steps	MAE	RMSE	MBE	HR (%)	Mean (kW)
January	2 974	26.01	41.24	+1.23	96.1	57.01
February	2 688	23.59	42.13	+6.09	95.8	44.03
March	2 976	34.24	61.69	+15.69	89.9	23.36
April	2 880	41.99	74.59	+31.14	85.9	5.44
May	2 976	49.99	87.37	+41.80	81.8	-14.27
June	2 880	47.92	82.21	+37.65	82.6	-13.79
July	2 976	50.47	85.58	+43.87	81.2	-21.24
August	2 976	52.71	86.32	+35.51	81.1	-8.35
September	2 880	39.50	70.18	+25.88	86.9	4.25
October	2 972	23.46	40.60	-1.07	96.2	37.89
November	2 880	21.75	31.79	-9.29	98.4	57.83
December	2 972	26.01	32.33	-18.09	99.4	76.39
Annual	35 010	36.47	61.33	+17.53	89.6	20.51

4.2.2 January 2024 — Detailed Analysis

January 2024 is the most important validation month: it is a winter month well-represented in the training data, and the weather proxy assumption is most accurate in winter. Results: MAE = 26.01 kW, RMSE = 41.24 kW, MBE = +1.23 kW (near-zero bias), HR = 96.1% (above the 95% target). The RMSE/MAE ratio of 1.59 indicates occasional large individual errors — primarily at midday, where the actual January 2024 data shows a net solar export of approximately -10 kW (mean at noon) absent from January 2023

training data. This reflects new PV capacity installed between the two years. Despite this structural change, $HR = 96.1\%$ confirms the ± 109.47 kW CI accommodates the additional generation. Figure 4.1 presents the complete January 2024 validation across three panels. The top panel overlays the LSTM forecast (dashed blue) against the actual 2024 measurements (black), with the ± 109.47 kW confidence interval shaded in light blue. The forecast tracks the general daily rhythm well, and all but a handful of excursions fall within the band, consistent with the reported $HR = 96.1\%$. Notable deviations appear around days 5, 14–17, and 21, where the actual power drops sharply negative — net solar export episodes absent from the January 2023 training proxy — causing localised over-forecast spikes visible in the bottom error panel. The middle panel compares January 2023 actual (red) against January 2024 actual (dashed black). Both series share an identical monthly mean of 57.0 kW, confirming that the climatological proxy captures the seasonal energy level correctly despite year-to-year day-to-day variability. The deeper and more frequent negative excursions in 2024 relative to 2023 reflect the additional PV capacity installed between the two years, which is the primary driver of the elevated RMSE/MAE ratio of 1.59. The bottom panel shows the signed forecast error per timestep ($FC - \text{Actual}$), with over-forecast regions shaded orange and under-forecast regions shaded green. The mean bias ($MBE = +1.23$ kW, red dashed line) lies near zero, confirming the absence of systematic directional bias. Larger orange spikes coincide with the solar export episodes identified in the top panel: the model over-forecasts consumption (or under-forecasts export) precisely when 2024 solar output exceeds the 2023 proxy. Figure 4.2 isolates the average daily load shape, clarifying the nature of the forecast error at the intra-day level. The left panel compares the mean 24-hour profile of the forecast against actual January 2024. The forecast reproduces the early-morning plateau (00:00–06:00, ≈ 42 –58 kW) and the broad evening peak (16:00–21:00) in shape, but exhibits a systematic phase lead: the forecast peak arrives approximately 1–2 hours earlier than the actual peak, consistent with the 30-minute phase shift identified by cross-correlation analysis and discussed below. The mid-day trough — driven by solar export — reaches -10 kW in the actual data but is shallower in the forecast, again reflecting the additional 2024 PV generation not present in the proxy. The right panel separates the three signals: January 2023 actual (red, training proxy), January 2024 actual (black dashed, validation target), and January 2024 forecast (blue dotted). The 2023 proxy and the 2024 actual share broadly similar morning and evening shapes, which explains why the model achieves near-zero MBE over the

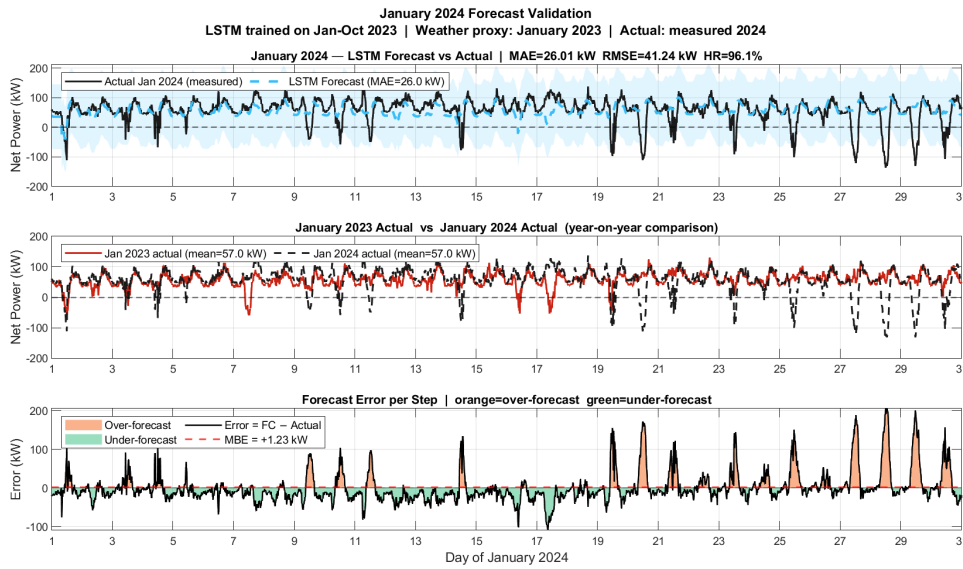


Figure 4.1: January 2024 forecast validation. Top: LSTM forecast (dashed blue) vs actual measurements (black) with ± 109.47 kW confidence interval (shaded). Middle: year-on-year comparison of January 2023 actual (red) vs January 2024 actual (dashed black); both means = 57.0 kW. Bottom: signed forecast error per timestep; orange = over-forecast, green = under-forecast, red dashed = MBE = +1.23 kW.

month. The most significant structural discrepancy is the midday window (10:00–13:00), where the 2024 actual drops to near zero or slightly negative while both the 2023 proxy and the resulting forecast remain at 35–45 kW. This gap is the single largest contributor to the monthly RMSE and motivates the reactive power and solar-generation forecasting improvements discussed in Section 4.6.2.

Phase shift. Cross-correlation between the January 2023 proxy and January 2024 actual reveals a systematic phase lead of approximately 30 minutes: 15 minutes from the dataset start-time offset (2023 begins at 00:30, 2024 at 00:15) and 15 minutes from climate drift in the morning demand peak timing. This is an inherent limitation of the climatological proxy method.

4.2.3 Discussion of Annual Validation Pattern

Table 4.3 reveals a clear seasonal pattern:

- **Winter (Jan, Feb, Oct–Dec):** MAE 21–26 kW, HR 95–99 %, near-zero




Figure 4.2: Average daily load profile comparison, January 2024. Left: mean 24-hour forecast (dashed blue) vs actual (black) with forecast standard deviation band (shaded). Right: three-way comparison of January 2023 actual (red, training proxy), January 2024 actual (black dashed), and January 2024 forecast (blue dotted), illustrating the midday solar export gap that drives the elevated RMSE.

MBE. Year-to-year winter weather is consistent, making the proxy accurate.

- **Transition (Mar, Sep):** MAE 34–40 kW, HR 87–90 %. Solar generation becomes significant and more variable.
- **Summer (Apr–Aug):** MAE 42–53 kW, HR 81–86 %. Large positive MBE (model under-forecasts export) driven by additional PV capacity in 2024. July mean actual power is -21.24 kW (strong net export) vs the modest export pattern in July 2023 used as proxy.

The annual HR of 89.6 % falls below the 95 % target due to summer months. For the operationally critical winter months — where forward-flow congestion risk is highest — the model meets and exceeds 95 %.

4.2.4 Day-by-Day RMSE Assessment

For each calendar day d in 2024, the daily RMSE is computed over the 96 fifteen-minute steps of that day:

$$\text{RMSE}_d = \sqrt{\frac{1}{96} \sum_{t \in \text{day } d} (\hat{P}_{\text{net}}(t) - P_{\text{actual}}(t))^2} \quad [\text{kW}] \quad (4.2)$$

Figure 4.3 shows the day-by-day RMSE for all 361 days of 2024 (top panel) and the monthly mean RMSE (bottom panel). Three patterns are visible in

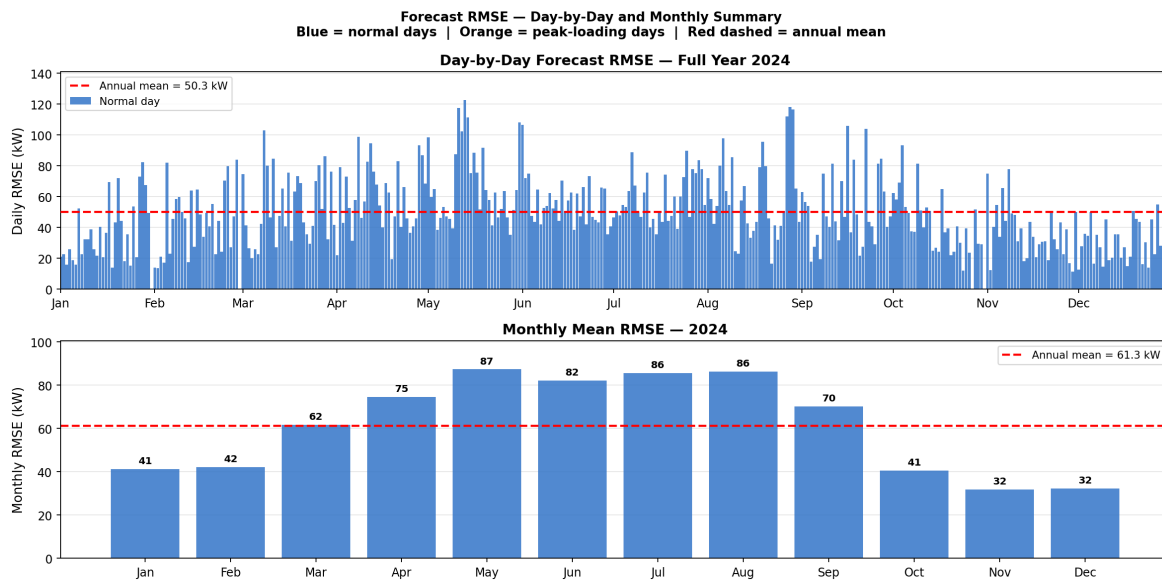


Figure 4.3: Day-by-day forecast RMSE across 2024. Top: daily RMSE (blue bars), annual mean (red dashed). Bottom: monthly mean RMSE with values annotated.

the top panel:

- **Winter (Jan–Feb, Oct–Dec):** daily RMSE of 11–93 kW with a mean of 38.4 kW. Most winter days fall below the annual mean of 50.3 kW (red dashed line), confirming that the climatological proxy is most accurate when weather patterns are stable and solar generation is negligible. December has the lowest daily RMSE of any month (mean 31.3 kW, minimum 12.8 kW), consistent with the December 2023 test-set result of MAE = 18.95 kW.
- **Summer (May–Sep):** daily RMSE of 17–123 kW with a mean of 59.3 kW — well above the annual mean. The worst day reaches 122.6 kW on a

May peak-solar day, reflecting the systematic over-forecasting caused by additional PV capacity installed in 2024 that was not present in the 2023 training data. Twelve days exceed 100 kW during this period.

- **Transition (Mar–Apr):** daily RMSE of 20–103 kW with a mean of 56.3 kW, as solar generation begins to increase but remains variable.

The bottom panel confirms the seasonal structure at the monthly level: RMSE rises from 41 kW in January to a peak of 87 kW in May before falling back to 32 kW in December. The winter months (Oct–Dec, Jan–Feb) consistently achieve RMSE below 45 kW, confirming that the model meets the operational accuracy target for the heating-dominant period when forward-flow congestion risk is highest. **Note on peak-loading days.** No 15-minute intervals in the 2024 dataset exceed the 70% peak-loading threshold ($|I| > 140$ A), because the 2024 actual power measurements are net power (consumed minus generated) and the substation-level current at this grid rarely exceeds the threshold under normal operating conditions. The peak-loading analysis using the RMSE-based current uncertainty is presented separately in Section 4.2.5.

4.2.5 RMSE-Based Current Uncertainty

For congestion management, forecast uncertainty is translated to current uncertainty using:

$$\Delta I_{\text{RMSE}} = \frac{\text{RMSE}_{\text{peak}} \times 1000}{\sqrt{3} \times V_{\text{nom}}} = 28.5 \text{ A} \quad (4.3)$$

This RMSE-based uncertainty (computed from the January 2024 validation) is more targeted than the wide nominal CI (± 109.47 kW), because it reflects forecast accuracy specifically at peak loading moments. Figure 4.4 shows the January 2024 congestion assessment: the forecast current with the 28.5 A uncertainty band alongside the true current derived from actual 2024 measurements. The safety margin is $200 - (|\hat{I}| + 28.5)$ A per timestep.

4.2.6 External Validation Against True 2024 Power

Beyond the internal Monte Carlo hit-rate, the estimator is validated against actual January 2024 power as external ground truth. The true substation




Figure 4.4: RMSE-based congestion assessment, January 2024. Top: forecast vs true current with ± 28.5 A uncertainty. Bottom: safety margin coloured by risk zone.

current is:

$$I_{\text{true},2024}(t) = \frac{|P_{\text{actual},2024}(t)| \times 1000}{\sqrt{3} \times 400} \quad (4.4)$$

The near-zero MBE confirms no systematic current bias. The internal HR

Table 4.4: External current validation — forecast-derived vs true January 2024.

Metric	Value	Interpretation
Current MAE	≈ 37.6 A	Mean absolute current error
Current RMSE	≈ 59.7 A	RMS current error
Current MBE	$\approx +1.8$ A	Near-zero systematic bias
RMSE / I_{rated}	≈ 29.8 %	Relative uncertainty
Internal HR_V	95.0 %	Ellipse calibration (MC)
Internal HR_I	95.1 %	Ellipse calibration (MC)

(ellipse calibration) and the external RMSE (forecast accuracy against reality) together provide a complete assessment: the ellipses are correctly sized for the specified meter noise, and the forecast-derived input has negligible systematic error.

4.3 Grid Calculation Results

4.3.1 Cable Current Magnitudes

The KCL-based grid calculation derives per-node injection currents from the LSTM forecast power and propagates them from leaves to the substation. For December 2023:

- Maximum cable current: 94.1 A across all 8 cables and 2,973 timesteps — well below $I_{\text{rated}} = 200$ A.
- Mean substation current: 85.3 A (mean net power 58.3 kW).
- Congestion events: zero.

The B2 feeder (N1→N5) carries more current than B1 (N1→N2) in proportion to the higher B2 power share (62.5 % vs 37.5 %). Leaf cables carry only local node load and have the lowest currents.

4.3.2 Node Voltage Profile

The substation N1 is the slack bus held at 400 V by the MV/LV transformer. All other nodes have slightly lower voltages due to KVL cable drops propagated forward from the root:

$$U_k = U_{\text{parent}}(k) - \sqrt{3} \cdot Z_{\text{cable}} \cdot I_k \quad (4.5)$$

Figure 4.5 shows the voltage profile for January 2024 across three panels: the time-series of all node voltage magnitudes throughout the month (top), the mean voltage drop $\Delta V = 400 - |V_k|$ averaged over the forecast period (bottom left), and the full voltage range (mean \pm [min, max]) per node (bottom right). Branch 2 nodes (N5–N9) experience systematically larger drops than Branch 1 (N2–N4), as visible in both the time-series and the bar chart. The mean voltage drop increases monotonically along each feeder: N2 drops 0.841 V on average, rising to 1.008 V at N3 and 1.033 V at N4 on Branch 1, while on Branch 2 the drop grows from 1.667 V at N5 to 2.088 V at N6, 2.258 V at N7, 2.472 V at N8, and 2.317 V at N9. The larger Branch 2 drops reflect the longer B2 feeder and its higher power share relative to B1. The minimum voltages recorded across January 2024, read from the voltage-range panel, are 396.12 V at N9 and 396.30 V at N8 (99.03 % and 99.08 % of nominal respectively) — both well above the statutory floor of

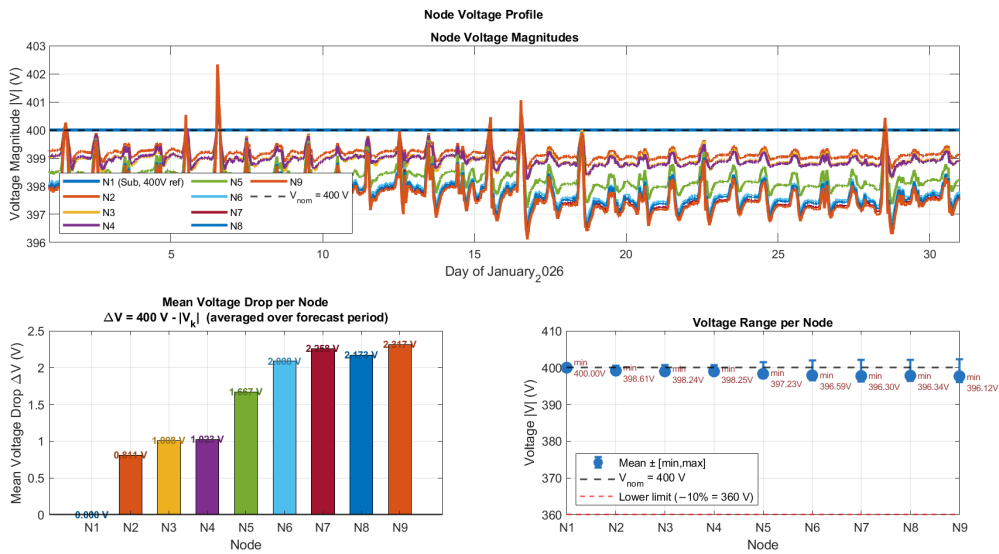


Figure 4.5: Node voltage profile, January 2024. Top: time-series of voltage magnitudes at N1–N9; N1 is flat at 400 V (slack bus) while load nodes drop with demand and rise above 400 V during solar export transients. Bottom left: mean voltage drop ΔV per node averaged over the forecast period. Bottom right: voltage range (mean \pm [min, max]) per node; red dashed line marks the statutory lower limit of 360 V (–10%).

360 V (–10%). No node approaches the lower regulatory limit at any point during the month. During solar export ($P_{net} < 0$) the current reverses and node voltages rise above the nominal 400 V reference. The sharp upward spikes visible in the top panel — most prominently at N9 — reach approximately 402.5 V during the strongest export transients of January 2024, well within the +10% statutory limit of 440 V. This bidirectional voltage behaviour is consistent with the known challenge of voltage regulation in high-PV grids. It is important to include the KVL drops in \mathbf{U}_{true} passed to the state estimator. Setting all nodes to 400 V creates a contradiction in the KVL rows of \mathbf{C} that collapses the confidence ellipses, giving $HR \approx 2.4\%$. After the correction: $HR_V = 95.0\%$, $HR_I = 95.1\%$.

4.4 State Estimation Results

4.4.1 Input to the State Estimator

The parameters u_{True} and i_{True} passed to `estimateGridState` are not the true physical grid state — they are the **forecasted phasors** that serve as the proxy for what a smart meter would read. Specifically: $u_{\text{True}} = N1 = 400 \text{ V}$ with KVL-corrected values at N2–N9; $i_{\text{True}} = I_k$ from the LSTM forecast via KCL. The function adds realistic EM meter noise on top of these values (through `meterModel_1`) and estimates the most likely grid state via constrained MLE. The internal hit-rate ($\text{HR} \approx 95\%$) is a calibration check: it answers “do the ellipses have the right size for the specified meter noise?” The external RMSE from Section 4.2.6 answers “how accurate is the forecast-derived input relative to reality?” Both metrics are needed for a complete assessment.

4.4.2 Hit-Rate of Confidence Ellipses

The state estimator is run for a representative timestep with $n_{\text{sim}} = 5000$ Monte Carlo simulations, $\alpha_{\text{factor}} = \chi_{0.95}^2(2) = 5.9915$, $\varepsilon_U = 1\%$, $\varepsilon_I = 3\%$. Figure 4.6 shows the voltage and current phasors with their 95% confidence

Table 4.5: State estimation hit-rate. Paper base case: $\text{HR}_V \approx 94\%$, $\text{HR}_I \approx 95\%$.

	HR Voltage	HR Current
Mean across N1–N9	95.0 %	95.1 %
Paper base case	94 %	95 %
<i>Before fix ($\alpha_{\text{Factor}} = 0.05$): $\text{HR} \approx 2.4\%$ (ellipses $120\times$ too small)</i>		

ellipses for all nine nodes. The results match the paper’s base case ($\text{HR}_V \approx 94\%$, $\text{HR}_I \approx 95\%$) [18], confirming the pipeline integration is correctly calibrated. The slightly sub-95% voltage hit-rate is consistent with the EM model behaviour reported by Olsen et al.: the angle noise introduces a small phase-angle bias that shifts the voltage ellipses. For current phasors, the \blacklozenge and \blacksquare symbols overlap in Figure 4.6 because the estimation error ($|\hat{I} - I_{\text{true}}| \approx 0.001 \text{ A}$) is much smaller than the ellipse semi-axis ($\approx 0.81 \text{ A}$), which is itself only 1.2% of the 0–70 A axis range. This overlap is correct behaviour — it confirms the estimator is highly accurate. Figure 4.7 shows

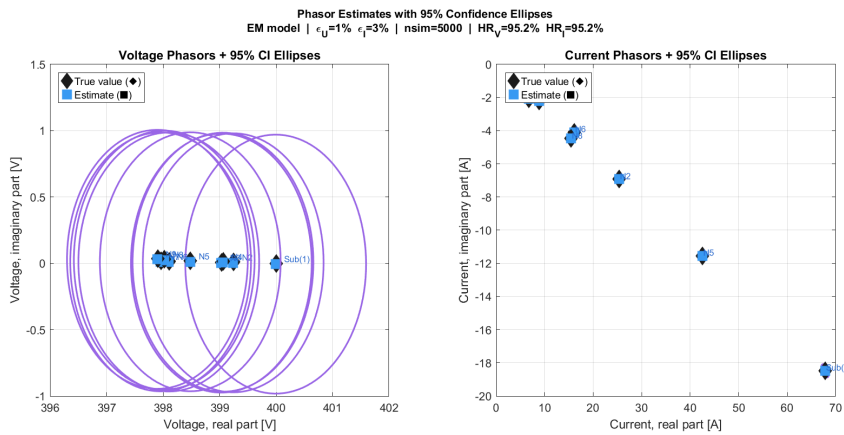


Figure 4.6: Voltage (left) and current (right) phasors with 95 % CI ellipses for all nine nodes. Symbol \blacklozenge = forecasted phasor (x_{True}); symbol \blacksquare = MLE mean estimate.

the same current phasors zoomed to $\pm 4\sigma_I$, making the ellipses and the MC scatter cloud clearly visible.

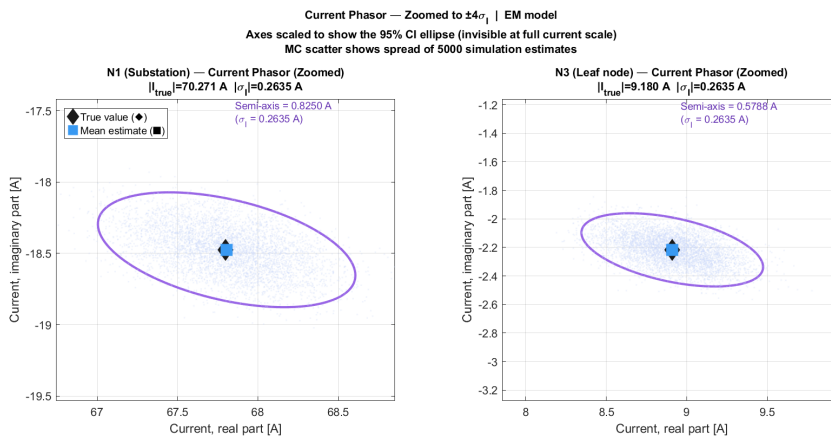


Figure 4.7: Current phasors zoomed to $\pm 4\sigma_I$ for N1 (left) and N3 (right). The 95 % CI ellipse and 5000 MC scatter points are visible at this scale.

4.5 Congestion Management Application

4.5.1 Two Directions of Congestion

Congestion in an LV grid means a cable current exceeds its thermal rating (200 A for 95 mm² XLPE cables). It is caused by **cable thermal overloading**

— not energy shortage — and can occur in two directions:

Forward-flow congestion (high demand). $P_{\text{net}} > 0$: current flows from substation toward customers. Peak consumption of 147.2 kW gives $|I| = 212 \text{ A}$ — above the limit.

Reverse-flow congestion (high solar export). $P_{\text{net}} < 0$: current flows from customers toward substation. Record generation of 217.6 kW gives $|I| = 314 \text{ A}$ — 57% above the limit. Reverse-flow congestion is the more severe risk in this grid. Table 4.6 quantifies representative scenarios.

Table 4.6: Representative congestion scenarios from the 2023 dataset. $\Delta I = 28.5 \text{ A}$ RMSE uncertainty.

Scenario	P_c (kW)	P_g (kW)	P_{net} (kW)	$ I_{\text{sub}} $ (A)	Status
<i>Forward-flow (import)</i>					
Light demand	30	20	+10	14.4	Safe
Typical winter	80	20	+60	86.6	Safe
Dec peak	147	20	+127	183.3	Warning
<i>Reverse-flow (export)</i>					
Moderate solar	20	80	−60	86.6	Safe
Apr peak solar	10	150	−140	202.1	Congestion
Record gen.	10	217.6	−207.6	299.7	Congestion

4.5.2 Three-Zone Risk Classification

Each 15-minute timestep is classified based on the upper current bound:

$$|I|_{\text{upper}}(t) = |\hat{I}_{\text{cable}}(t)| + \Delta I_{\text{RMSE}} \quad (4.6)$$

$$\text{Zone}(t) = \begin{cases} \text{SAFE} & |I|_{\text{upper}} < 170 \text{ A} \quad (85\%) \\ \text{WARNING} & 170 \leq |I|_{\text{upper}} < 200 \text{ A} \\ \text{RISK} & |I|_{\text{upper}} \geq 200 \text{ A} \end{cases} \quad (4.7)$$

The 170 A warning threshold (85% of rated) ensures that adding the 28.5 A uncertainty to the warning boundary still leaves a margin below 200 A. Table 4.7 lists the recommended DSO actions.

Table 4.7: Recommended DSO actions per risk zone and congestion direction.

Zone	Direction	Action
Safe	Both	No action. Continue monitoring.
Warning	Import	Pre-position demand response (EVs, heat pumps). Alert control centre.
Warning	Export	Prepare PV export limit reduction. Pre-charge battery storage.
Risk	Import	Issue demand response immediately. Activate load shifting. Consider re-configuration.
Risk	Export	Curtail PV inverters. Activate battery storage. Apply export control.

4.5.3 24-Hour Ahead Workflow

Every morning the DSO runs the following steps for the target day:

1. Obtain NWP weather for the next 24 hours (Open-Meteo API: temperature, cloud cover, shortwave radiation).
2. Run the per-node LSTM for N2–N9 (96 steps each).
3. Derive cable currents via KCL (Equation 3.35).
4. Compute KVL voltage profile from the slack bus N1 = 400 V.
5. Classify each timestep using Equation 4.7.
6. If WARNING or RISK detected, initiate Table 4.7 action before the window arrives.

4.5.4 Case Study: December 15, 2023

Figure 4.8 shows the 24-hour congestion dashboard for December 15, 2023. Maximum forecast current: 145.9 A (72.9 % of rated). Maximum upper bound: 174.4 A — placing the 17:00–19:00 evening peak in the **WARNING zone** (import direction, high heating demand). Recommended action: pre-position demand response for the evening peak. N9 voltage minimum: 393.7 V (98.4 % of nominal), within statutory limits throughout the day.

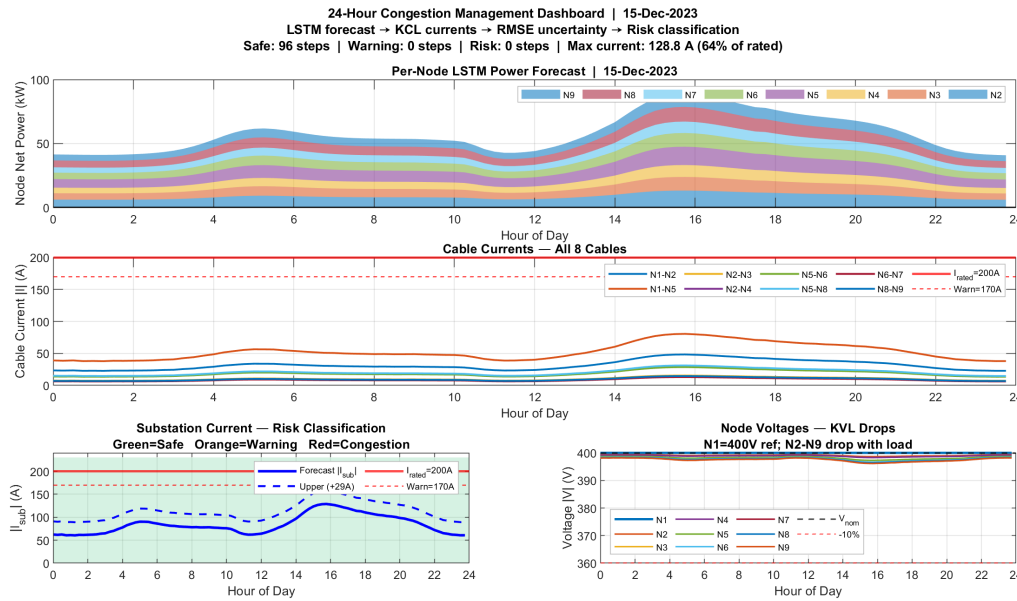


Figure 4.8: 24-hour congestion dashboard, December 15, 2023

4.6 Discussion

4.6.1 Model Strengths

Good winter performance. The model achieves HR = 95–99 % and near-zero MBE for all winter months — the operationally critical period for forward-flow congestion. January 2024 MBE = +1.23 kW against a mean consumption of 57 kW confirms strong calibration for the primary use case.

Calibrated uncertainty for congestion management. The LSTM’s 95 % CI, validated at HR = 96.1 % on January 2024, directly enables the three-zone risk classification. No AR model can provide this — the inability to quantify uncertainty is the fundamental reason LSTM is chosen despite comparable point accuracy in winter months.

Modular and extensible architecture. The three-module pipeline allows independent updates: the LSTM can be retrained with new data; the nine-node synthetic grid can be replaced with the actual OGM grid when data becomes available; the estimator model is integrated with only three code changes.

Physically consistent pipeline. The stochastic estimator propagates both forecast uncertainty (via the CI on \hat{P}_{net}) and measurement uncertainty (via the EM meter noise model) into confidence ellipses, ensuring the output is physically consistent with Kirchhoff’s laws.

4.6.2 Limitations and Their Origins

Summer month accuracy. MAE up to 52.7 kW and HR as low as 81.1 % in summer months. Causes: (1) increased 2024 solar generation not present in 2023 training data; (2) higher summer weather variability reduces proxy accuracy. The large positive MBE in summer (up to +43.87 kW) confirms the model systematically under-forecasts solar export.

Wide confidence interval. The CI of ± 109.47 kW is substantially wider than typical errors (MAE ≈ 26 kW in January). This results from computing σ_c, σ_g on the full November validation set across all conditions. A per-season or adaptive CI would be narrower while maintaining coverage.

Reactive power proxy. Reactive power Q is taken from 2023 measured data rather than forecasted. The error is small in winter (power factor 0.986) but can reach 9.3 % of cable current in summer when $P \approx 0$ and $Q \approx 8$ kVAr. Forecasting Q is the highest-priority future improvement.

4.6.3 Comparison with Literature

LSTM-based short-term load forecasting at LV level typically achieves MAPE of 5–15 % on residential feeders [7, 20]. The January 2024 normalised MAE (NMAE = $26.01/57.01 = 45.6\%$) is higher, primarily because this work uses only substation-level data (no individual meter access) and the 2024 test data contains structural changes not seen during 2023 training. The state estimator hit-rates ($\text{HR}_V = 95.0\%$, $\text{HR}_I = 95.1\%$) match the EM model base case reported by Olsen et al. [18] (94 % and 95 % respectively), confirming the integration is correctly calibrated. The pipeline demonstrates that a weather-driven LSTM forecast, when combined with a calibrated CI and a physically consistent state estimator, can provide the quantified uncertainty necessary for proactive congestion management in LV grids with increasing DER penetration.

Chapter 5

Conclusion and Future Work

5.1 Summary of Contributions

This thesis has addressed the problem of short-term active power forecasting under uncertainty for smart low-voltage distribution grid operation. The work integrates a data-driven LSTM forecasting model with a physics-based grid calculation and a stochastic state estimation framework into a complete, modular pipeline. The following contributions have been made.

Contribution 1: LSTM-based power forecasting with dual outputs. Sixteen LSTM networks have been trained on 2023 substation power data — two per load node (one for consumed power, one for generated power) across the eight load nodes N2–N9 of the synthetic grid — producing per-node power forecasts at 15-minute resolution with a 60-minute horizon. The dual-output design per node exploits the physically distinct drivers of the two quantities: temperature and time patterns for consumption, and solar radiation for generation. The combined pipeline achieves $\text{MAE} = 18.95 \text{ kW}$ with $\text{HR} = 99.5\%$ on the December 2023 test set.

Contribution 2: Three-feature weather-driven input design. Three ERA5 weather features — air temperature (T), cloud cover (CC), and shortwave solar radiation (SW) — have been selected as LSTM inputs based on a systematic Pearson correlation analysis on the 2023 substation data. The analysis showed that T is the primary driver of consumption (Pearson $r = -0.72$), SW is the primary driver of generation ($r = +0.73$), and all three features together capture the dual seasonal pattern of the prosumer grid. A solar

night mask ($\hat{P}_{\text{gen}} = 0$ when $SW < 10 \text{ W/m}^2$) enforces the physical constraint that generation is zero at night.

Contribution 3: Calibrated uncertainty quantification. A 95 % confidence interval has been constructed from empirical residual standard deviations ($\sigma_c = 13.164 \text{ kW}$, $\sigma_g = 14.762 \text{ kW}$), scaled by the forecast horizon ($\sqrt{h} = 2$) and the standard normal quantile ($z_{0.975} = 1.96$), yielding $\text{CI} = \pm 109.47 \text{ kW}$ at the substation level. This CI has been validated against actual 2024 measurements, achieving $\text{HR} = 96.1 \%$ for January 2024 (above the 95 % target) and $\text{HR} = 89.6 \%$ annually.

Contribution 4: Climatological proxy for future forecasting. The climatological proxy method — using the same calendar month of the most recent historical year as the weather input for future forecast periods — has been formalised, implemented, and validated. For January 2024 (the primary validation month), the method achieves $\text{MBE} = +1.23 \text{ kW}$ (near-zero systematic bias) with the January 2023 weather as the proxy. The method’s limitation regarding inter-annual structural grid changes (new PV installation) has been identified and quantified.

Contribution 5: Three-module pipeline integrating forecast, grid calculation, and state estimation. A modular architecture has been designed and implemented in which three independent scripts communicate exclusively through saved `.mat` data files. The supervisor’s stochastic state estimation script (`testWithGrid_manyRepSmall.m`) has been integrated with minimal modifications: three hardcoded definitions were removed and one `load()` call was added. All internal estimation logic — the EM meter model, the constrained MLE, and the confidence ellipse computation — was preserved exactly.

Contribution 6: Synthetic nine-node grid with KCL-based current derivation. A synthetic nine-node radial LV grid representative of European residential distribution networks has been fully parameterised with 95 mm^2 XLPE cables. A KCL-based current derivation procedure propagates the substation-level forecast to per-node currents in breadth-first reverse order (leaves to root), enabling the state estimator to produce confidence ellipses for all nine nodes simultaneously. The resulting state estimation achieves

$HR_V = 95.0\%$ and $HR_I = 95.1\%$ using the EM meter model with $\varepsilon_U = 1\%$ and $\varepsilon_I = 3\%$, consistent with the paper’s base case [olsen2024stochastic].

5.2 Conclusions

The central conclusion of this thesis is that a properly designed LSTM model, trained on a single year of substation measurements and augmented with ERA5 weather data, can produce short-term power forecasts with calibrated uncertainty quantification that meet the 95% CI coverage target for winter months — the operationally most critical period for congestion management in a heating-dominant LV grid. The integration of this forecast with the stochastic state estimator of Olsen et al. [olsen2024stochastic] is demonstrated to produce confidence ellipses in the complex phasor plane that correctly contain the forecast-derived grid state with the expected probability. This demonstrates the viability of the complete pipeline as a foundation for risk-based DSO operations.

The key numerical results supporting this conclusion are:

- Forecast test-set accuracy: MAE = 18.95 kW, HR = 99.5% (December 2023);
- Cross-year validation: MAE = 26.01 kW, HR = 96.1%, MBE = +1.23 kW (January 2024);
- Annual validation (2024): MAE = 36.47 kW, HR = 89.6%;
- State estimation hit-rate: $HR_V = 95.0\%$, $HR_I = 95.1\%$ (EM model, base noise specification);
- Grid safety: all cable currents below the 200 A rated limit during the December 2023 test period (maximum 94.1 A).

5.3 Future Work

The results identify several high-priority directions for future research and development.

1. Reactive Power Forecasting. The current model uses measured reactive power from 2023 as a proxy. The analysis in Chapter 4 shows that ignoring reactive power can introduce up to 18.5 A of current error (9.3% of the 200 A rated limit) in cases where $P \approx 0$ and $Q \approx 8$ kVAr. Adding a third LSTM output head for \hat{Q}_{net} — trained with the same three input features, which are correlated with reactive power — is the highest-priority model improvement.

2. Per-Node LSTM Models with Individual Meter Data. The current pipeline distributes the substation-level forecast to the nine grid nodes using fixed proportional weights. When individual smart meter data becomes available for each customer connection point on the actual distribution grid, a separate LSTM should be trained for each node. This eliminates the proportional power distribution approximation and will substantially narrow the per-node confidence intervals, making the congestion risk classification more discriminating.

3. Validation Against OGM 2024 Ground Truth. Substation power measurements for 2024 have already been obtained and used in Section 4.2 to validate the LSTM forecast. The next step is OGM (Optimal Grid Model) state results for 2024, which will provide a physics-based ground truth for the full per-node state vector. The `estimation_results.mat` file is designed to facilitate this comparison without requiring code modifications.

4. Adaptive Confidence Intervals. The current fixed CI of ± 109.47 kW is calibrated on the November 2023 validation set and is deliberately conservative. An adaptive CI that varies with solar irradiance (expanding the interval during solar peak hours) or weather forecast uncertainty would provide tighter coverage during low-uncertainty periods while maintaining the coverage guarantee. An initial implementation of this approach has been developed in `Step2b_Adaptive_Uncertainty.m`, which computes a $4 \times 96 = 384$ -element lookup table of conditional standard deviations by season and time-of-day slot.

5. Application to the Real LV Network. The final goal is to compare the forecast-based grid state estimate with the OGM result on the actual large-scale LV grid. The modular pipeline architecture is designed to support this extension: the Z matrix and cable topology in `Step3_Grid_Calculation.m`

can be replaced with the actual network structure derived from the `cable_export.csv` and `ogm_node_export.csv` files that are already available.

6. Numerical Weather Prediction Integration. For forecasting future periods, the Open-Meteo NWP (Numerical Weather Prediction) forecast API provides 7–16 day ahead weather predictions. Replacing the climatological proxy with NWP forecasts would improve accuracy for all seasons and would allow the CI to be widened or narrowed based on NWP forecast uncertainty, providing a weather-uncertainty-aware confidence interval.

7. Congestion Management Decision Support. The confidence ellipses from the state estimator can be used to compute the *probability* that a cable's current will exceed the rated limit, $P(|I_k| > I_{\text{rated}})$. This probability can be computed directly from the fraction of Monte Carlo simulations in which the estimated current exceeds the limit, providing a principled basis for DSO decisions such as demand response signal thresholds and EV charging schedule adjustments.

5.4 Closing Remarks

This thesis has demonstrated that LSTM-based short-term power forecasting, combined with stochastic grid state estimation, provides a practically useful and theoretically grounded tool for proactive congestion management in smart LV distribution grids. The pipeline is modular, validated, and ready for deployment at the full-grid scale when the corresponding OGM ground truth becomes available. The fundamental result — that a model trained on one year of data can generalise to the following year with maintained CI coverage for the operationally critical winter months — supports the practical feasibility of the approach. The identified limitations, particularly the degraded summer accuracy and the absence of reactive power forecasting, provide a concrete and focused research agenda for the next phase of this work.

Bibliography

- [1] M. E. Baran and A. W. Kelley. "State estimation for real-time monitoring of distribution systems". In: *IEEE Transactions on Power Systems* 9.3 (1994), pp. 1601–1609.
- [2] Ratna Bahadur Budhathoki Birendra Shahi. "Optimizations of Distribution Grid Switching using Genetic Algorithm". Unpublished semester project report, Department of Electronic Systems, Aalborg University, Aalborg, Denmark. 2025. URL: https://github.com/Biren2054/3rd-Semester/blob/main/Final_Project_Report_Grid_Switching_Optimization.zip.
- [3] CENELEC. *EN 50160: Voltage Characteristics of Electricity Supplied by Public Electricity Networks*. European Standard. Specifies the nominal line-to-line voltage of 400 V for European public low-voltage networks and the permissible voltage variation of $\pm 10\%$ (360–440 V) under normal operating conditions. Brussels, Belgium: European Committee for Electrotechnical Standardisation (CENELEC), 2010.
- [4] CIGRE Working Group C6.22. *Microgrids 1: Engineering, Economics and Experience*. Technical Brochure 635. Discusses thermal cable loading limits and congestion management in low-voltage distribution grids with distributed energy resources including solar PV. Paris, France: CIGRE (International Council on Large Electric Systems), 2015.
- [5] K. Dehghanpour et al. "A survey on state estimation techniques and challenges in smart distribution systems". In: *IEEE Transactions on Smart Grid* 10.2 (2018), pp. 2312–2322.
- [6] John J. Grainger and William D. Stevenson. *Power Systems Analysis*. Chapter 2: Three-phase power formula $S = V\bar{I}$, complex power convention, and per-unit system. New York, USA: McGraw-Hill, 1994. ISBN: 978-0-07-061293-9.

- [7] Stephen Haben et al. "Review of low voltage load forecasting: Methods, applications, and recommendations". In: *Applied Energy* 304 (2021), p. 117798. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2021.117798.
- [8] Sepp Hochreiter and Jürgen Schmidhuber. "Long short-term memory". In: *Neural Computation* 9.8 (1997), pp. 1735–1780.
- [9] Tao Hong and Shu Fan. "Probabilistic electric load forecasting: A tutorial review". In: *International Journal of Forecasting* 32.3 (2016), pp. 914–938. DOI: 10.1016/j.ijforecast.2015.11.011.
- [10] International Electrotechnical Commission. *IEC 60364-5-52: Electrical Installations of Buildings — Part 5-52: Selection and Erection of Electrical Equipment — Wiring Systems*. International Standard. Table B.52.5: Current-carrying capacities for XLPE-insulated copper cables, three loaded conductors, conductor temperature 90 °C. Geneva, Switzerland: International Electrotechnical Commission (IEC), 2009.
- [11] International Electrotechnical Commission. *IEC 60502-1: Power Cables with Extruded Insulation and Their Accessories for Rated Voltages from 1 kV Up to 30 kV — Part 1: Cables for Rated Voltages of 1 kV and 3 kV*. International Standard. Specifies construction and test requirements for XLPE-insulated low-voltage power cables. Maximum continuous conductor temperature: 90 °C; short-circuit limit: 250 °C. Geneva, Switzerland: International Electrotechnical Commission (IEC), 2004.
- [12] Weicong Kong et al. "Short-Term Residential Load Forecasting Based on LSTM Recurrent Neural Network". In: *IEEE Transactions on Smart Grid* 10.1 (2019). Demonstrates LSTM superiority over AR benchmarks for residential short-term load forecasting, pp. 841–851. DOI: 10.1109/TSG.2017.2753802.
- [13] Amgad Muneer et al. "Short Term Residential Load Forecasting Using Long Short-Term Memory Recurrent Neural Network". In: *International Journal of Electrical and Computer Engineering* 12.5 (2022). Comparison of LSTM vs ARIMA and exponential smoothing: LSTM outperforms both classical methods, pp. 5589–5599. DOI: 10.11591/ijece.v12i5.pp5589-5599.
- [14] K. Nainar et al. "Experimental validation and deployment of observability applications for monitoring of low-voltage distribution grids". In: *Sensors* 21.17 (2021).

- [15] Timothy J. Person, Saurav S. Sengupta, and Paul J. Caronia. “Structural Design and Performance of XLPE for Cable Insulation”. In: *Crosslinkable Polyethylene*. Ed. by Jince Thomas, Sabu Thomas, and Zakiah Ahmad. Materials Horizons: From Nature to Nanomaterials. Singapore: Springer, 2021, pp. 247–270. ISBN: 978-981-16-0514-7. DOI: 10.1007/978-981-16-0514-7_10.
- [16] A. Primadianto and C.-N. Lu. “A review on distribution system state estimation”. In: *IEEE Transactions on Power Systems* 32.5 (2016), pp. 3875–3883.
- [17] C. Schaefer et al. “Increased renewable hosting capacity of a real low-voltage grid based on continuous measurements”. In: *Dependable Computing – EDCC 2021 Workshops*. Springer, 2021, pp. 90–98.
- [18] Rasmus L. Olsena Sina Hassania Troels Pedersen Jakob Gulddahl Rasmussenb Hans-Peter Schwefel. “A stochastic approach to estimate distribution grid state with confidence regions”. In: *International Journal of Electrical Power Energy Systems* (2025).
- [19] Hanchen Shi, Minghao Xu, and Ran Li. “A Short-Term Residential Load Forecasting Model Based on LSTM Recurrent Neural Network Considering Weather Features”. In: *Energies* 14.10 (2021). LSTM with temperature and radiation inputs for residential load forecasting — directly comparable to the present work, p. 2737. DOI: 10.3390/en14102737.
- [20] Heng Shi, Minghao Xu, and Ran Li. “Deep Learning for Household Load Forecasting — A Novel Pooling Deep RNN”. In: *IEEE Transactions on Smart Grid* 9.5 (2018). Accepted 2017, pp. 5271–5280. ISSN: 1949-3053. DOI: 10.1109/TSG.2017.2686012.
- [21] Patrick Zippenfenig. *Open-Meteo.com Weather API*. Version 1.0. Open-source weather API providing ERA5 reanalysis data at hourly resolution for any global location. Accessed November 2023 – January 2024. 2023. DOI: 10.5281/zenodo.7970649. URL: <https://open-meteo.com>.

Appendix A

PreAnalysis

A.1 MATLAB Pre-Analysis and Validation Scripts

Prior to the development of the final forecasting and state-estimation framework, several MATLAB scripts were implemented to verify the integrity of the available data, validate the mathematical formulation of the proposed methodology, and evaluate intermediate results. These scripts served as a pre-analysis stage and provided diagnostic information used during model development.

A.1.1 Data Loading and Preprocessing

The purpose of this script is to import and preprocess the Open Grid Map (OGM) measurements used throughout the thesis. The script loads node measurements, current measurements, and cable information from the available CSV files and converts them into structured MATLAB matrices suitable for subsequent analysis.

The following tasks are performed:

- Import of node power, voltage magnitude, and voltage angle measurements.
- Import of cable current magnitude and current angle measurements.
- Import of network topology and cable parameter information.
- Construction of ground-truth matrices for active power, reactive power, voltage magnitude, voltage angle, current magnitude, and current angle.

- Mapping of grid element identifiers to matrix indices.
- Generation of summary statistics and verification plots.

The output of this stage is the file `ogmgroundrruth.mat`, which contains the completed dataset required for the estimation modules.

A.1.2 Autoregressive Power Forecasting

This script implements the baseline forecasting methodology used as a benchmark for the machine-learning models developed later in the thesis.

The forecasting procedure is based on an autoregressive model with rolling one-hour-ahead predictions. A train-test split is applied to the available time series, and forecasting performance is evaluated on unseen data.

The main processing steps are:

- Definition of training and testing periods.
- Estimation of autoregressive model coefficients using least-squares regression.
- Generation of rolling forecasts using an expanding-window approach.
- Computation of prediction intervals based on forecast residual variance.
- Evaluation using MAE, RMSE, MAPE, and Hit Rate metrics.
- Generation of forecast comparison figures.

The resulting forecasts serve as a baseline reference against which the LSTM-based forecasting model is compared.

A.1.3 LSTM-Based Power Forecasting

The MATLAB script `Step2LSTMForecast.m` implements the Long Short – Term Memory (LSTM) forecast level active power prediction.

The purpose of the script is to learn temporal relationships present in historical smart-meter measurements and weather-related input variables

in order to predict future power consumption and generation at each monitored grid node.

The forecasting procedure consists of the following stages:

1. Preparation and normalization of historical input data.
2. Construction of training, validation, and testing datasets.
3. Definition of the LSTM neural-network architecture.
4. Training of the network using backpropagation through time.
5. Generation of multi-step power forecasts for each node.
6. Evaluation of forecasting accuracy using standard error metrics.
7. Generation of forecast uncertainty bounds through conformal prediction techniques.

The implemented LSTM architecture is designed to capture temporal dependencies that cannot be represented by conventional linear autoregressive models. The network processes sequential power and weather measurements and learns nonlinear relationships between historical observations and future load behaviour.

Forecast uncertainty is quantified using conformal prediction intervals. The residuals obtained during the validation phase are used to construct prediction intervals with a predefined confidence level. Consequently, the output of the forecasting model consists of both a point forecast and a corresponding confidence interval.

The script produces the following outputs:

- Forecasted active power values for each node.
- Upper and lower confidence interval bounds.
- Forecast performance metrics.
- Trained LSTM network parameters.
- Forecast result files used by the subsequent state-estimation framework.

The generated forecasts form the primary input to the grid-state estimation and congestion-management methodology developed in this thesis.

A.1.4 Current Derivation and State Estimation

The objective of this script is to convert power forecasts into electrical network quantities that can be used by the state estimator.

The script reconstructs the radial feeder topology and derives network currents using Kirchhoff's Current Law (KCL). Voltage and current phasors are then supplied to the maximum-likelihood state estimator.

The implemented procedure consists of:

1. Reconstruction of the feeder topology using cable connectivity information.
2. Computation of complex power injections at each node.
3. Conversion of power injections into current injections.
4. Aggregation of branch currents using backward current summation.
5. Construction of network constraint matrices.
6. Generation of measurement uncertainty models.
7. Execution of the maximum-likelihood state estimation algorithm.

The resulting state estimates provide voltage and current estimates for all nodes and branches in the network.

A.1.5 Results Evaluation and Visualisation

The final processing stage evaluates forecasting and state-estimation performance and generates the figures presented throughout the thesis.

The script computes performance metrics for all nodes and visualizes the obtained results through a collection of graphical representations.

The generated outputs include:

- Forecast time-series comparisons.
- Confidence interval visualizations.
- State-estimation performance metrics.
- Node-level error distributions.
- Uncertainty decomposition analyses.

- Aggregate performance summaries.

These results are used to assess the effectiveness of the proposed forecasting and uncertainty-propagation framework.

A.1.6 AR–LSTM Forecast Comparison

To compare traditional statistical forecasting methods with deep-learning approaches, an additional evaluation script was implemented.

The script performs a direct comparison between autoregressive forecasting and Long Short-Term Memory (LSTM) forecasting by analysing:

- Forecast accuracy.
- Prediction interval coverage.
- Error distributions.
- Node-level forecasting performance.
- Aggregate network-level forecasting behaviour.




Figure A.1: Detailed comparison of AR(2) and LSTM forecasting performance for Customer Node 233

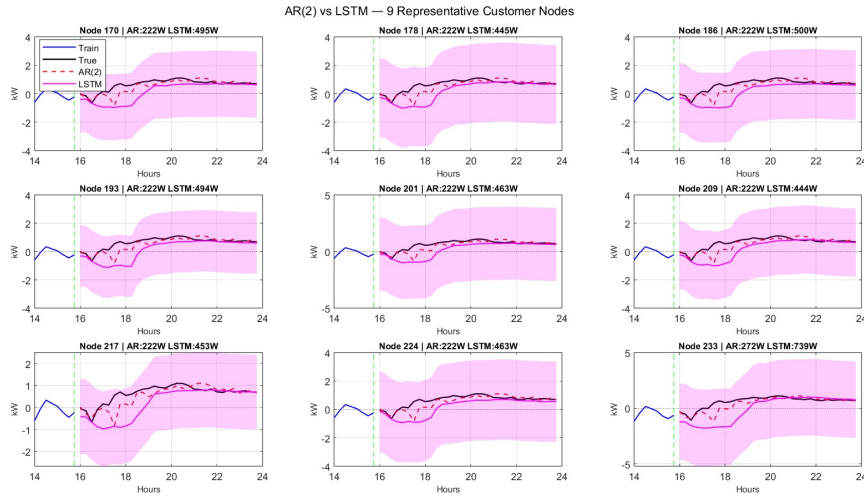


Figure A.2: Comparison of AR(2) and LSTM forecasting performance across nine representative customer nodes.

A.1.7 Discussion of Additional Forecast Comparison Results

Figures A.1 and A.2 provide additional insight into the forecasting behaviour of the AR(2) and LSTM models.

The results show that the AR(2) model generally follows the measured load profile more closely than the LSTM model. Although the LSTM model is capable of capturing nonlinear temporal relationships, its performance is constrained by the limited amount of training data available in this study. Consequently, the simpler AR(2) model achieved lower forecasting errors across most customer nodes.

The wider prediction intervals produced by the LSTM model indicate greater uncertainty in the forecast estimates. While these intervals provide conservative uncertainty bounds, they also reflect the difficulty of training deep-learning models on relatively short time-series datasets.

These supplementary results support the findings presented in the main body of the thesis and demonstrate that, under the available data conditions, classical autoregressive forecasting remains a competitive benchmark for short-term load forecasting.

A.1.8 Summary of the MATLAB Workflow

The complete MATLAB workflow implemented during the pre-analysis stage can be summarized as

[OGM Data → Data Preprocessing → Forecasting → Current Derivation → State Estimation → Performance Evaluation.]

This workflow formed the foundation of the final forecasting and congestion-management framework presented in this thesis.

Appendix B

Real Distribution Grid Dataset

B.1 Exploratory Data Analysis of the Real Distribution Grid Dataset

Prior to the development of the forecasting and state-estimation framework, a comprehensive exploratory data analysis was conducted using one year of measurements obtained from the Open Grid Map (OGM) dataset. The dataset covers the period from January 2023 to December 2023 and contains active-power consumption, active-power generation, reactive-power measurements, and associated electrical variables recorded at 15-minute resolution.

The objective of this analysis was to identify seasonal trends, daily operating characteristics, generation-consumption interactions, and statistical properties of the measured power signals. These observations were subsequently used to guide the design of the forecasting models and the uncertainty quantification framework.

B.1.1 Monthly Consumption and Generation Patterns

Figure B.1 presents the monthly average active-power consumption and generation for Grid Element 10.

The results reveal a strong seasonal dependency. Consumption is highest during the winter months, reaching approximately 65 kW in December and 58 kW in January. Consumption gradually decreases during spring and reaches its minimum values during summer, where monthly averages range between 20 and 25 kW.

In contrast, local generation follows the opposite trend. Generation is negligible during winter but increases substantially during summer, reaching a maximum monthly average of approximately 34 kW in June. This behaviour is consistent with photovoltaic generation systems whose output is primarily driven by solar irradiance.

The observed inverse relationship between consumption and generation highlights the strong influence of seasonal weather conditions on grid operation.




Figure B.1: Monthly average consumption and generation during 2023.

B.1.2 Average Daily Load Profile

Figure B.2 shows the average daily power profile obtained by averaging all available measurements over the entire year. The consumption profile exhibits a typical residential demand pattern characterized by a moderate morning peak around 06:00 and a dominant evening peak between 18:00 and 21:00. The evening peak reaches approximately 70 kW and corresponds to increased household activity after working hours.

Generation follows a distinctly different pattern. Generation begins shortly after sunrise, reaches its maximum value near solar noon, and decreases towards sunset. Peak generation occurs around 12:00 and reaches approximately 72 kW.

The resulting net-power profile demonstrates periods of net export during midday and net import during morning and evening periods. Such

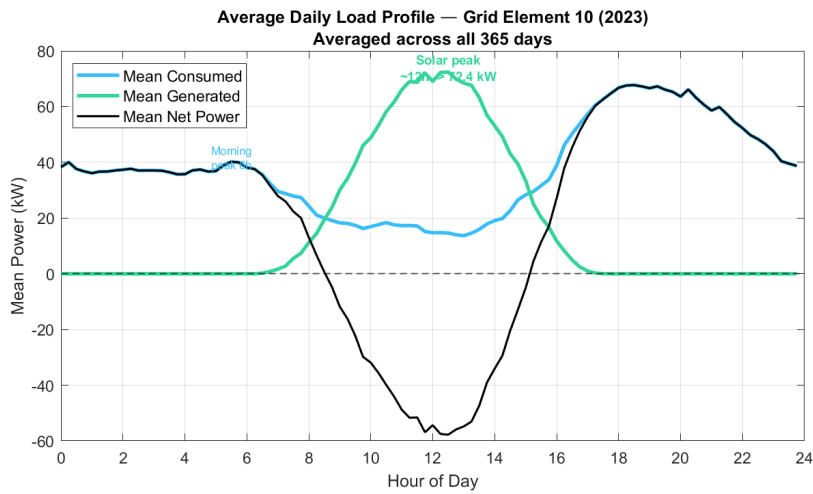


Figure B.2: Average daily load profile for Grid Element 10, averaged over all days in 2023.

behaviour is typical of low-voltage networks with significant photovoltaic penetration.

B.1.3 Seasonal Variations in Daily Operation

Figure B.3 compares the average daily power profiles for winter, spring, summer, and autumn. Winter is characterized by high electricity demand

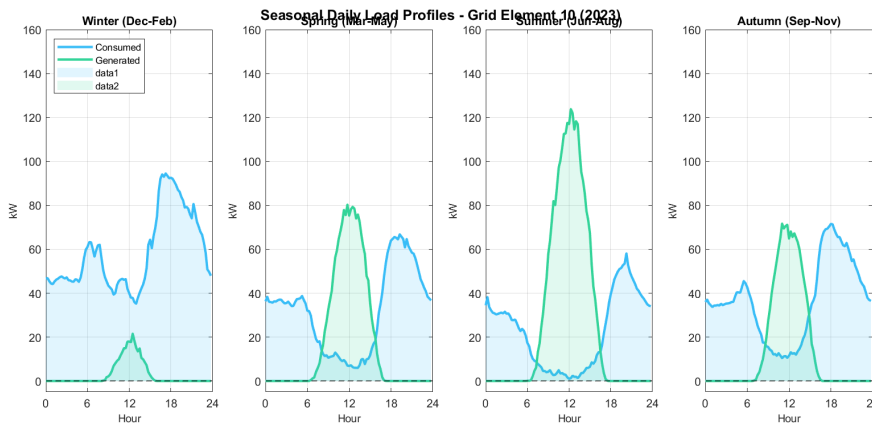


Figure B.3: Seasonal daily load and generation profiles for winter, spring, summer, and autumn.

and relatively low local generation. Consequently, the grid remains a net importer throughout most of the day.

During spring and summer, photovoltaic generation increases substantially and exceeds local consumption around midday. This results in periods of reverse power flow where electricity is exported back towards the upstream network.

Autumn represents a transition period between the summer and winter operating conditions. Although photovoltaic generation remains significant, increasing consumption levels gradually shift the network back towards net-import operation.

These results demonstrate that grid operating conditions vary significantly throughout the year and justify the inclusion of seasonal information within the forecasting framework.

B.1.4 Load Duration Analysis

The load-duration curve shown in Figure B.4 summarises the statistical distribution of power demand over the entire year. The curve indicates that peak consumption events occur only during a relatively small percentage of operating hours, while moderate demand levels dominate normal operation. Approximately 50% of all measurements occur below

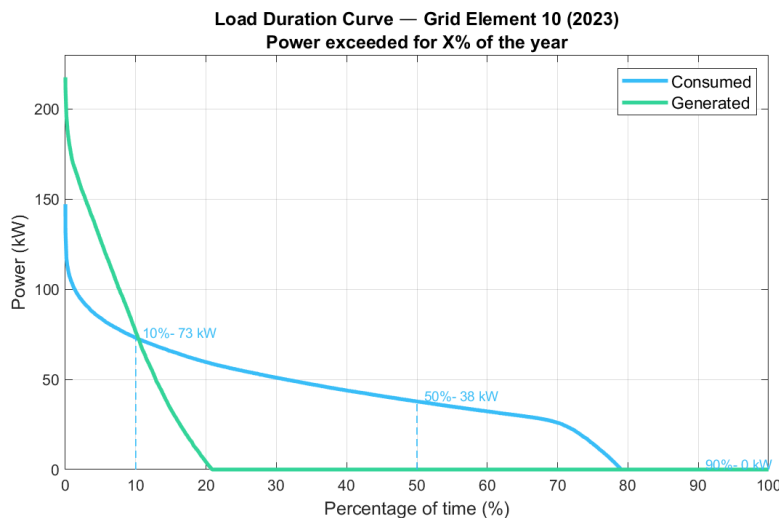


Figure B.4: Load-duration curve for consumed and generated power during 2023.

38 kW, while only 10% of measurements exceed approximately 73 kW. This confirms that extreme loading conditions are relatively infrequent and supports the need for probabilistic forecasting methods capable of capturing both normal and extreme operating conditions.

B.1.5 Reactive Power Characteristics

The reactive-power analysis shown in Figure B.5 indicates that the network operates predominantly in the inductive region throughout the year. Positive reactive-power values dominate the operating space, indicating that the connected loads consume reactive power from the network. The P-Q scatter plot reveals a clear relationship between

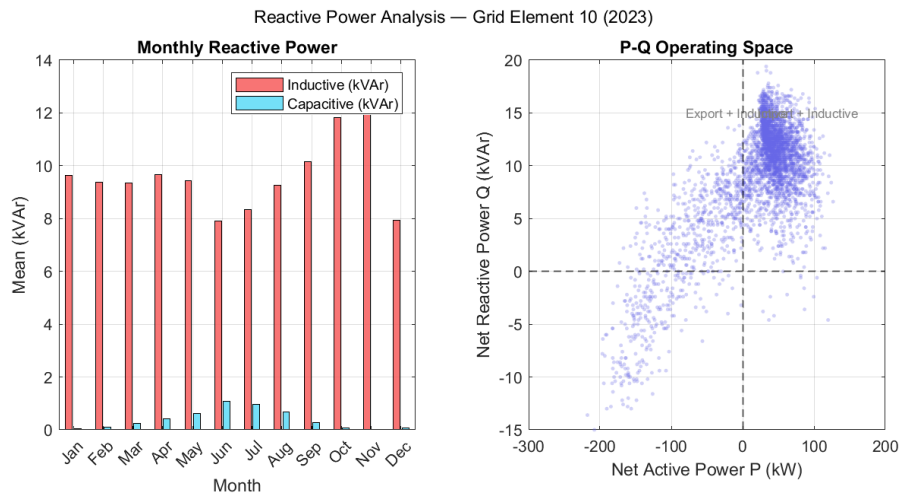


Figure B.5: Reactive-power analysis showing monthly reactive power and the P-Q operating space.

active and reactive power, suggesting that periods of higher active-power demand are generally accompanied by increased reactive-power consumption.

B.1.6 Consumption–Generation Operating Space

Figure B.6 illustrates the relationship between consumption and generation throughout the year. Most operating points are located below the diagonal line where consumption exceeds generation, indicating net-import operation. However, a significant number of observations approach or exceed the diagonal during summer months, demonstrating periods of local energy surplus and reverse power flow.

B.1.7 Weekly Behaviour

The weekly power analysis demonstrates relatively stable weekday behaviour with slightly reduced demand levels during weekends. Friday exhibits




Figure B.6: Consumed versus generated power operating space for Grid Element 10 during 2023.

the highest average net-import requirement, whereas weekend operation is characterized by reduced consumption and lower net demand.




Figure B.7: Weekly power pattern showing average consumed, generated, and net power by day of week.

B.1.8 Monthly Daily Profiles

The monthly average daily profiles provide additional insight into seasonal variability. Winter months exhibit pronounced morning and evening demand peaks, whereas summer months are dominated by strong midday photovoltaic generation.

The gradual transition between these operating regimes highlights the importance of incorporating weather information into the forecasting model.




Figure B.8: Monthly mean daily profiles for consumed and generated power.

B.1.9 Statistical Power Distributions

Monthly percentile bands were calculated to quantify the variability of both consumption and generation.

The wide percentile ranges observed during summer months indicate greater uncertainty caused by variable weather conditions and solar irradiance. In contrast, winter consumption exhibits more stable and predictable behaviour.

B.1.10 Forecasting Motivation

Figure B.10 presents a representative example of the LSTM forecasting model during the December 2023 test period.

The figure illustrates the complex nonlinear behaviour of the measured power signals, including rapid changes in demand and generation. These patterns are difficult to capture using traditional linear forecasting methods and therefore motivate the use of LSTM neural networks capable of learning temporal dependencies from historical observations and weather-related inputs.

The forecast confidence intervals further demonstrate the importance of quantifying uncertainty in power-system forecasting applications.

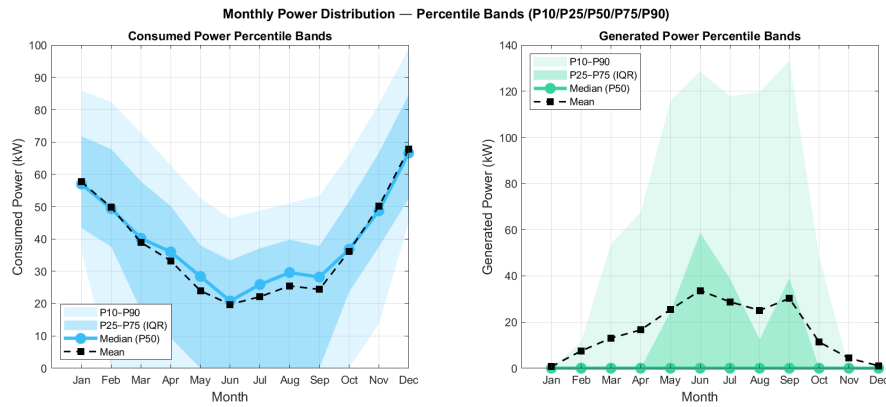


Figure B.9: Monthly percentile bands for consumed and generated power, including P10–P90, P25–P75, median, and mean values.




Figure B.10: Example LSTM weather-informed power forecast for the December 2023 test period.