ACADEMIA International Journal for Social Sciences
Volume 5, Issue 3(s9), 2026 ISSN-L (Online): 3006-6638

Demand-Side Management of EV Integration with Lesco Grid Station

Mudasir Rafique
mudassarrafique737@gmail.com
Department of Electrical Engineering, The Superior University Lahore, Pakistan

Dr. Manzoor Ellahi
manzoor.ellahi@superior.edu.pk
Department of Electrical Engineering, The Superior University Lahore, Pakistan

Wasif Arshad
arshadwasif2022@gmail.com
Department of Electrical Engineering, The Superior University Lahore, Pakistan

Zafrullah Khan
zafarullah.khan@superior.edu.pk
Department of Electrical Engineering, The Superior University Lahore, Pakistan

Corresponding Author: Dr. Manzoor Ellahi manzoor.ellahi@superior.edu.pk

Received: 17-01-2026 **Revised:** 01-02-2026 **Accepted:** 15-02-2026 **Published:** 02-03-2026

ABSTRACT

With an estimated 500,000 EVs on Pakistani roads by 2030, Pakistan's National Electric Vehicle Policy (NEVP) 2019 aims for 30% of all new vehicle sales to be electric. This transformation is an acute threat to the infrastructure of Lahore Electric Supply Company (LESCO) which is already operating 146 grid stations having 10/13 MVA (132/11 kV) ONAN/ONAF power transformers with 80-90% loading of their ONAN capacity on peak summer evenings. At 30% penetration, the uncoordinated Level-2 EV charging causes the transformer to be loaded to 15.57 MVA (119.8% of its 13MVA ONAF transformer rating) for 2.5 hours per day, bringing the hot-spot temperature above the IEEE C57.91 rated limit of 110°C by 37.1°C and increasing the daily insulation loss-of-life by about 28 times compared with no EV charging. In this paper, a modular AI-Driven Demand-Side Management (AI-DSM) framework is proposed that addresses this crisis using only software - no hardware changes required. The framework consists of three successive steps: (1) an unsupervised identification of EV owners using k-means and DBSCAN clustering on load profiles of smart meters; (2) a 24 h-ahead transformer loading forecast based on a stacked Long Short-Term Memory (LSTM) neural network; (3) a quadratic-programme based charging-schedule optimiser that fills the overnight load valleys while taking into account the thermal capacity of transformers and the departure constraints of EV owners. The integrated framework has been validated against a 12-month simulation based on a calibrated LESCO Annual Report for 2022–23, and yields: EV identification F1 = 1.000 with zero false positives; load forecasting R^2 = 0.974, MAE = 0.168 MVA (53.8% improvement over seasonal naive baseline); peak load reduction 36.1% (transformer loading restored to 76.5% of ONAF rating); winding hot-spot held at 86.1°C (versus no-EV baseline of 85.2°C); and 100% user charging satisfaction — all without capital investment in hardware.

Keywords: *Electric Vehicle Integration, AI-Driven Demand-Side Management, LESCO, LSTM Load Forecasting, K-means Clustering, Power Transformer Ageing, IEEE C57.91, Quadratic Programming, Non-Intrusive Load Monitoring*

BACKGROUND

The global transportation sector is undergoing an electrification revolution. By 2030, the International Energy Agency projects a global EV fleet of 350 million, with EVs accounting for approximately 60% of all new car sales [1]. Pakistan has recognised this imperative: the National Electric Vehicle Policy (NEVP) 2019 targets 30% of new vehicle sales as electric and approximately 500,000 EVs on Pakistani roads by 2030 [2]. This transition is essential for decarbonisation and energy security, but it also creates unprecedented challenges for an electricity distribution network designed and built decades before the concept of residential EV charging existed.

LESCO supplies electricity to the Lahore and adjacent districts which have more than 6.5 million customers. LESCO has 146 grid stations with 10/13 MVA (132/11 kV) ONAN/ONAF power transformers, each supplying around 47000 distribution transformers to the residential and commercial consumers [3]. One of the defining features of LESCO's network is its load profile which is seasonal in nature: during the peak summer evenings, when the temperature creeps to 45°C or above and the load is primarily dominated by air-conditioning demand, the transformers at the Grid stations are already operating at 80-90% of their 10 MVA ONAN natural-cooling capacity, thus leaving the transformers with little headroom for any further demand.

The 10/13 MVA grid-station power transformer is the strategic asset of concern. Each unit serves the aggregated demand of thousands of consumers; its accelerated ageing or failure under prolonged overload has network-wide consequences, with replacement costs and procurement lead times far exceeding those of any single distribution transformer. The Arrhenius reaction-rate model codified in IEEE Std C57.91-2011 [4] governs the life of transformer winding insulation: for every 6–8°C that the winding hot-spot temperature exceeds the rated threshold of 110°C, insulation life is approximately halved. Even brief, repeated overloads accumulate irreversible thermal damage.

The analyses have consistently revealed that EV charging will shift electricity demand to the peak hours when residential demand is already the highest. Muratori (2018) showed, by using a feeder simulation in a residential area, that EVs tend to cluster in space as a consequence of social influence and peer effects, with a penetration of 50-70% in some neighbourhoods despite a low penetration at the system level [5]. Hilshey et al. (2013) and Gong et al. (2012) quantified this probabilistically - uncoordinated charging can shorten a transformer's effective service life from 40 years to less than 10 [6,7]. Shafiee et al., (2013) reported 30-40% increase in peak transformers loading at 15-20% EV penetration, which is above thermal design limits [8]. One of the complicating issues is what is called in this paper the “blindness problem” — EV purchases are not communicated to utilities, making it impossible to plan ahead, and making the transformer overloads seem random.

Motivation

There are two types of solutions that have been suggested to integrate EVs into the grid. The cost of hardware solutions such as transformer upgrade, dedicated EV feeders, advanced metering infrastructure is PKR 50-100 million per grid station, or PKR 7-15 billion for all 146 grid stations of LESCO [3]. This is not possible for a developing country to utility, which has competing infrastructure needs. Simple software interventions like Time-of-Use (ToU) tariffs lower the average peak demand but are unable to cater for the dynamic and spatially heterogeneous adoption of EVs across LESCO's service territory, nor can they ensure transformer safety in high penetration areas [10].

Thus an intelligent, software-based DSM system is needed that works within LESCO's current AMI system (without extra sensors and communication devices), and can deliver actionable operational intelligence: (i) identify EV-owning households based on consumption data, without being provided with labelled training data; (ii) predict transformer loading 24 hours in advance with accurate enough results to plan for it; and (iii) calculate the optimal charging windows to prevent transformer overload while satisfying user mobility requirements. These 3 capabilities correspond to the 3 modules of the AI-DSM proposed in this paper.

LITERATURE REVIEW

EV Impact on Distribution Transformers

Clement-Nyns et al. (2010) demonstrated that uncoordinated PHEV charging causes substantial power losses and voltage deviations beyond the capacity of conventional protection schemes [11]. The distribution transformer emerges as the principal bottleneck: its cellulose-paper-and-mineral-oil insulation degrades exponentially with temperature per the Arrhenius model in IEEE Std C57.91-2011 [4], and a transformer operated consistently 12°C above its rated hot-spot temperature has an effective insulation life roughly one-fifth of its design value. Hilshey et al. (2013) showed uncoordinated charging can raise the equivalent ageing factor by orders of magnitude; Gong et al. (2012) demonstrated that a few EVs charging simultaneously in hot-spot conditions readily exceed the critical threshold [6,7]. Muratori (2018) added the spatial dimension: EV adoption clustering means specific transformers can face local penetrations of 50–70%, creating failure hotspots invisible to system-level planning [5]. Most of this literature is grounded in North American or European smart-grid environments with adequate below-feeder visibility and less heavily loaded assets, leaving a critical contextual gap for developing-grid utilities such as LESCO where below-feeder visibility is limited and assets are already heavily loaded.

EV Load Identification — NILM

The concept of Non-Intrusive Load Monitoring (NILM) to identify appliances from their steady state real and reactive power signature was first introduced by Hart (1992) [9]. Modern deep learning methods like the Sequence-to-Point architecture of Zhang et al. (2020) have high disaggregation accuracy, however they demand a large amount of labelled training sets, which are not available at LESCO. [12] Parson et al. (2012) used the probabilistic disaggregation based on HMM, but these methods have difficulty in coping with high power loads that overlap [13]. An alternative to this is unsupervised clustering, which provides a label free approach. K-means is used extensively in practice and is a quick algorithm, yet it assumes spherical clusters. DBSCAN can find arbitrary shape clusters, and can explicitly identify outliers, which are attractive for the high power, time concentrated EV signature. Cominola et al. (2017) surveyed unsupervised methods and pointed out that they are not as accurate as supervised methods, but they do offer the scalability required for the deployment in real world where labels are limited [14]. There has been no study that has thoroughly compared the DBSCAN and k-means algorithms for the particular application of residential EV detection specifically in a developing-country environment without ground truth labels.

Short-Term Load Forecasting

Although the classical ARIMA and SARIMA models work well for smooth aggregated regional demands, they are not suitable to model the nonlinear and bursty nature of demands influenced by EVs (Raza and Khosravi, 2015) [15]. The standard for short-term load forecasting has been the Long Short-Term Memory (LSTM) architecture by Hochreiter and Schmidhuber (1997) [16] which can learn time series with long range temporal dependency with the help of gated memory cells. Kong et al. (2019) also confirmed the superiority of LSTM for household-level prediction, which offers a bursty pattern of residential demands, which is much more difficult to be captured by feed-forward networks, 35–40%

reduction of MAPE than traditional baselines [17]. Marino et al. (2016) have presented stacked LSTM models for energy forecasting for buildings, achieving additional improvement. No previous studies have specifically validated an LSTM-based transformer loading forecaster for LESCO's grid-station level in the climate of Pakistan where temperature-load correlation is very high and air-conditioning is the main residential load in summer.

Research Gap

The literature has verified that EVs are a very serious thermal risk for distribution transformers, LSTM is

the state-of-the-art load forecasting tool and that the unsupervised clustering can be used to detect EV loads without labels. What is missing is an end-to-end, hardware-free framework that is capable of combining all three capabilities, and tested on a real utility in a developing country that has limited resources and is particularly interested in being able to monitor thermal lifetime and the condition of the insulation of power transformers at the grid station, according to IEEE C57.91. This paper fills this gap.

Problem Statement

The following interconnected conditions define the problem:

- On a peak summer evening, LESCO's 10/13 MVA grid-station transformers are operating at about 91% of their ONAN rating (~9.1 MVA) with a diversified residential load, thus providing very limited capacity to support any extra EVs without exceeding their rated thermal envelope.

- For 2.5 hours of peak loading at 15.57 MVA – 119.8% of the ONAF rated 13MVA – due to uncoordinated post-work charging at 30% residential EV penetration (1200 Level-2 chargers at 7 kW per grid-station catchment of ~4000 households), the peak winding hot-spot temperature is 122.3°C, 12.3°C above the IEEE C57.91 rated limit of 110°C.

- Thermal damage is exponential: under uncoordinated charging, daily loss-of-life for thermal damage is about 28 times no-EV baseline and is a direct reduction to the service life of the transformer.

- Under uncoordinated charging, the thermal limit is exceeded at as little as 22% EV penetration, which Pakistan's EV adoption curve will hit in the next decade in high-income urban areas.

- LESCO does not have a means to determine who the EV owners are; thus, without the ability to identify described in this paper, targeted DSM interventions are impossible.

- Financing for providing hardware support at PKR 50 – 100 million per grid station is not feasible at a large scale at LESCO's 146 grid stations and 47,000 distribution transformers.

Proposed AI-DSM Framework

Framework Architecture

The AI-DSM framework is a sequential pipeline of three modules, where the output from one module is passed as input to the next. Module 1 (EV Identification) identifies households with EVs and those that do not, based on the annual smart-meter load profile without relying on any labelled training data, by unsupervised clustering. Module 2 (Load Forecasting) takes the historical transformer loading, current weather forecast and calendar features and creates a 24-hour ahead load forecast at 15-minute intervals.

The calculation of an overnight charging schedule is done in module 3 (Charging Optimisation) using the LSTM forecast as the basis load, and the roster of EVs households as the list of controllable loads, to minimise peak loading whilst respecting the following constraints: in transformer capacity, user energy needs, and departure-time. Each module is testable separately yet is a complete pipeline from raw meter data to actionable scheduling decisions. The conceptual architecture is shown in figure 1.

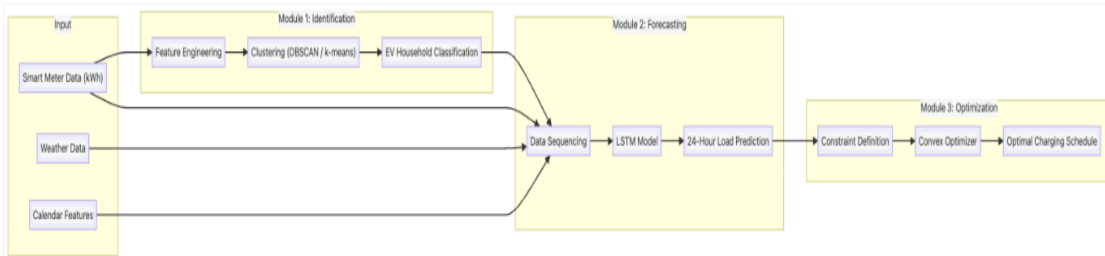

Figure 1: AI-DSM Framework — three-module sequential pipeline from smart-meter data to actionable EV charging schedules. [Source: Dissertation Figure 3.1]

Network Scope

The major scope is LESCO's 10/13 MVA, 132/11 kV grid-station power transformer, ONAN 10MVA and ONAF 13MVA. The catchment model is about 4,000 residential connections. Further, the framework is transferable down the network hierarchy as a secondary validation case is performed at the 200 kVA, 11 kV/400 V distribution transformer level (80 households) (Table 1).

Asset	Rating	Voltage	Connections	Role in Study
Grid-Station Power Transformer	10/13 MVA (ONAN/ONAF)	132/11 kV	~4,000 households	Primary analysis target
Distribution Transformer	200 kVA	11 kV/400 V	80 households	Secondary validation

Table 1: Network Hierarchy and Scope

System Formulation and Methodology

A. Data Sources and Preprocessing

The pipeline has three streams of data. To create the residential load profiles, a bottom-up approach is used by combining appliance duty cycle data and the appliance ownership statistics in Lahore households and then aggregating to feeder and grid-station level calibrated as per consumption statistics in LESCO Annual Report 2022-23 [3]. The resulting profiles are similar to LESCO's typical evening peak period (18:00-22:00) at ~9.1 MVA and are consistent with the reported per-consumer annual consumption. Weather information including the ambient temperature (in °C), relative humidity (in %) and solar irradiance is collected from Open-Meteo API Lahore. Since the summer air-conditioning load in Lahore is significantly related to ambient temperature (at > 45°C on peak days), temperature is considered the most important exogenous predictor. EV charging profiles are assumed to be Level 2 AC charging sessions with a constant charging rate of 7 kW (32 A, 220 V, single-phase) and with a charging duration depending on the battery size and State of Charge (SOC) (Table 2). All data are at 15 minutes resolution for a 12-month period, split temporally into 70/15/15% training, validation, and test sets, respectively, to avoid temporal data leakage (Tables 3 and 4). Preprocessing includes linear interpolation for gaps < 1 hour, forward fill for gaps > 1 hour, 3σ clipping on outliers, Min-Max normalisation for inputs into the neural network and one-hot encoding on day-of-week.

Parameter	Value	Rationale
Charger type	Level 2 AC (single-phase)	Most common residential mode internationally and in Pakistan
Rated power	7 kW (32 A, 220 V)	Standard Level-2 rating
Mean daily energy / EV	20 kWh	Typical battery fill for daily commute distance
Mean arrival time	18:55 (post-commute)	Probabilistic model; coincides with evening peak
Mean departure time	06:49 (next morning)	~12 h scheduling window available
Data resolution	15 minutes	Consistent with AMI rollout specification in Pakistan
Observation period	12 months	Covers all seasonal variations in Lahore

Table 2: EV Charging Session Parameters and Dataset Configuration

Dataset Component	Detail
Residential households	80 profiles; 24 EV-owning (30% penetration), 56 baseline
Grid-station load record	12 months, 15-min resolution (35,040 time steps)
EV identification ground truth	Known by construction — used only for algorithm evaluation
Exogenous features	Temperature (°C), humidity (%), solar irradiance, day-of-week (one-hot), holiday flag

Table 3: Dataset Summary

Partition	Share	Period	Purpose
Training	70%	Months 1–8.4	LSTM weight fitting; clustering parameter selection
Validation	15%	Months 8.4–10.2	Early stopping; LSTM hyperparameter tuning
Test	15%	Months 10.2–12	Unbiased out-of-sample performance evaluation

Table 4: Chronological Data Partitioning Strategy (70/15/15%)

B. Module 1 — EV User Identification

The objective is to classify the ~4,000 residential connections in the grid-station catchment as EV-owning or non-EV-owning using only smart-meter consumption data — without labelled training data or customer self-reporting.

Feature Engineering

Ten statistical features are computed per household from the full-year daily load profile (Table 5). Feature selection is guided by the physical characteristics of Level 2 charging: a 7 kW charger activated in the post-commute evening window (18:00–24:00) creates a step-change in household power consumption which has no parallel among common Pakistani residential loads. An inverter air-conditioner — the next largest appliance — cycles by compressor modulation and rarely exceeds 2.5 kW per unit. Features therefore capture three aspects of the charging signature: (i) temporal concentration in the evening window; (ii) magnitude of the power step-change; and (iii) the resulting shape distortion of the daily load curve. All features are z-score normalised to zero mean and unit variance before clustering.

Feature	Definition	Physical Rationale
Evening mean (kW)	Mean load 18:00–24:00	Post-commute charging elevates evening average
Evening std (kW)	Standard deviation of load 18:00–24:00	Charging events increase within-window variance
Evening energy share (%)	% of daily energy consumed in 18:00–24:00	Temporal concentration of EV load
Daily maximum (kW)	Peak load over 24 hours	7 kW charger amplifies household daily peak
Daily range (kW)	Max minus min load	Amplitude of the step-change signature
Peak-to-trough ratio	Daily max / daily min load	Ratio amplified by the 7 kW

		step-change
Max ramp rate (kW/15 min)	Largest load increase in any 15-min interval	7 kW switch-on has no residential analogue
Load factor	Mean load / peak load	Inversely related to EV-induced peakiness
Daily energy (kWh)	Total daily consumption	EV households consume ~20 kWh more per day
Overnight mean (kW)	Mean load 00:00–06:00	Non-discriminative control feature (validated)

Table 5: Ten Features Extracted for EV Identification Clustering

K-means Implementation

K-means is configured with k = 2 (EV / non-EV), k-means++ initialisation (reducing sensitivity to initial centroid placement), Euclidean distance metric, convergence tolerance 10^{-4}, and a maximum of 300 iterations. Since k is known as a priori from the binary classification objective, k-means has a single tuneable hyperparameter; its O(n) per-household computational complexity makes it well-suited for automated screening across LESCO's ~600,000 metered connections. Implementation uses Python Scikit-learning.

DBSCAN Implementation

To use DBSCAN (Density-Based Spatial Clustering of Applications with Noise), the parameters MinPts = 2 × 10 = 20 (twice the feature dimensionality, as suggested by Ester et al., 1996) and ε = 4.368 (from the elbow point of the sorted k-th nearest-neighbour distance graph, with k = MinPts) are applied. DBSCAN is added as a theoretically more powerful comparator, which is also capable of identifying clusters of arbitrary shape and of labelling noise points (outliers) which is appropriate to the high power, time-concentrated EV charging signature. It is implemented in Python Scikit-learn based on Euclidean distance.

Evaluation Metrics

The standard classification metrics are calculated: Precision, Recall, F1-Score, and Confusion Matrix, due to the fact that ground truth labels are available since the EV households are known from the data generation procedure. The Silhouette Score measures internal cluster separation, independent of labels (range −1 to +1; > 0.7 means good structure). Robustness is determined by adding gaussian noise (σ = 0 to 1.0) to the feature matrix to represent meter inaccuracy, artefacts from imputation, and unusual household behaviors.

C. Module 2 — LSTM Load Forecasting

The goal is to forecast the apparent-power loading of the 10/13 MVA grid-station power transformer over the next 24 hours (96-time steps at 15-minute resolution) for the charging optimiser to use as a baseline load curve.

Model Architecture

A stacked Long Short-Term Memory (LSTM) network is used (Table 6). At each time step t, the input vector comprises: the lagged transformer load $L(t-1, \ldots, t-96)$ providing 24 hours of history; ambient temperature $\theta(t)$ (°C); relative humidity $H(t)$ (%); day-of-week one-hot encoding (7 binary features); and a public-holiday flag (binary). Temperature is the dominant exogenous predictor due to the strong air-conditioning–temperature coupling in Lahore. This 96-step input window is processed through two stacked LSTM layers (64 and 32 hidden units respectively), each followed by Dropout (rate = 0.2) for regularisation, and a Dense output layer that generates the next 96 quarter-hourly load values. Stacking two LSTM layers allows the first layer to learn short-term temporal patterns (daily cycles, sharp EV ramps) while the second captures longer-range dependencies (seasonal temperature trends, weekly consumption rhythms).

Hyperparameter	Value	Rationale
LSTM Layer 1 hidden units	64	Captures short-term temporal patterns
LSTM Layer 2 hidden units	32	Captures longer-range dependencies at reduced dimension
Dropout rate	0.2 (after each LSTM layer)	Regularisation; prevents overfitting on 12-month dataset
Input history window	96 steps (24 hours)	Captures full preceding day's load cycle
Forecast horizon	96 steps (24 hours ahead)	Required for day-ahead scheduling
Optimiser	Adam (lr = 0.001)	Adaptive learning rate; robust to sparse gradients
Loss function	Mean Squared Error (MSE)	Penalises large forecast errors quadratically
Early stopping patience	10 epochs on validation loss	Prevents overfitting without fixed epoch count
Framework	TensorFlow / Keras (Python)	GPU-compatible; standard for deep learning

Table 6: LSTM Model Architecture and Hyperparameters

Forecasting Evaluation Metrics

Metric	Formula	Interpretation
MAE	Mean|y_pred − y_true|	Average absolute error in MVA; interpretable magnitude
RMSE	Sqrt(Mean(y_pred−y_true)²)	Penalises large errors quadratically; critical for safety
MAPE	Mean|y_pred−y_true|/y_true × 100%	Scale-free relative error (%)
R^2	1 − SS_res/SS_tot	Proportion of variance explained (1.0 = perfect)

Table 7: Forecasting Evaluation Metrics

D. Module 3 — Charging Optimisation

Three charging strategies are compared to quantify the value of each level of coordination (Table 8). The AI-DSM strategy is formulated as a Quadratic Programme (QP).

Scenario	Description	Data Required	Coordination Level
A — Uncoordinated	EVs charge immediately on arrival at full rated power; no utility control	None	None (worst case)
B — FCFS	Charging queued by arrival time; suspended reactively if aggregate load approaches ONAF rating	Real-time load monitoring	Reactive
C — AI-DSM Optimised	QP schedule computed from 24 h LSTM forecast; proactively fills overnight valley within 90% ONAF ceiling	LSTM day-ahead forecast + EV roster	Predictive (full optimisation)

Table 8: Charging Scenario Definitions

Optimisation Problem Formulation

Let P_base(t) be the LSTM-predicted baseline load at time step t, P_{c,i} the rated charging power of EV i, and x_i(t) ∈ [0, 1] the fractional charging power of EV i at time t. The objective function minimises the sum of squared aggregate load over the 24-hour scheduling horizon — a quadratic penalty that naturally suppresses peaks relative to valleys (load flattening / valley filling):

$$\min \sum_{t=1}^{T} \left[P_{base}(t) + \sum_{i=1}^{N} P_{c,i} \cdot x_i(t) \right]^2$$

where T = 96 (15-min steps) and N is the number of identified EV households. The squared formulation creates an incentive to spread charging load evenly rather than concentrating it, ensuring valley-filling behaviour without requiring an explicit valley-identification step.

Subject to four constraint classes:

- Capacity constraint: total load ≤ 0.90 × 13 MVA = 11.7 MVA at all time steps. The 10% safety margin (1.3 MVA buffer) absorbs forecast uncertainty — the LSTM ±2σ error band of ±0.46 MVA is well within this margin.
- Energy constraint: each EV i receives its full required energy E_{i,req} before departure, ensuring 100% user mobility satisfaction.
- Charging window: x_i(t) = 0 for all t outside [arrival_i, departure_i], enforcing plug-in and departure times.
- Variable bounds: 0 ≤ x_i(t) ≤ 1 for all i, t (fractional charging control).

A summary of the complete optimisation problem is shown in Table 9. Implementation with Python CVXPY (convex optimisation library); solving time is negligible, for the scale of N ≈ 1200 EVs.

Element	Specification
Decision variable x_i(t)	[0,1] fractional charging power for EV i at time step t
Objective	Minimise sum of squared aggregate load — promotes peak reduction and valley filling
Capacity constraint	Total load ≤ 11.7 MVA (90% of 13 MVA ONAF rating) at all T = 96 time steps
Energy constraint	Full required energy delivered to each of N EVs before their departure time
Charging window	x_i(t) = 0 outside [arrival_i, departure_i] for each EV i

Scheduling horizon T	96-time steps × 15 min = 24 hours
Scale N (30% penetration)	~1,200 EVs (grid-station level) / 24 EVs (distribution-transformer level)
Solver	CVXPY quadratic programme (Python)

Table 9: Charging Optimisation Problem Summary

E. Transformer Thermal Model — IEEE C57.91-2011

Transformer loading profiles are converted to winding hot-spot temperature and insulation loss-of-life using the IEEE C57.91-2011 dynamic thermal model [4]. The model describes two coupled first-order thermal systems: the oil thermal mass and the winding thermal mass, each with a distinct time constant. This is critical for accurate ageing assessment: because the oil has a much larger thermal mass than the winding, the hot-spot peak lags and is amplified relative to the load surge — a static approximation would significantly underestimate peak hot-spot temperature during transient overloads.

Top-Oil Temperature Model

The ultimate top-oil temperature rise over ambient at a given loading K (per-unit, referenced to 13 MVA ONAF) is:

$$\Delta\theta_{TO,ult} = \Delta\theta_{TO,R} \cdot \left[\frac{K^2 R + 1}{R + 1}\right]^n$$

where R is the ratio of rated load loss to no-load loss, and n is the oil viscosity exponent. The top-oil temperature evolves exponentially towards this ultimate value with oil time constant τ_TO:

$$\Delta\theta_{TO}(t) = \Delta\theta_{TO}(t-\Delta t) + [\Delta\theta_{TO,ult} - \Delta\theta_{TO}(t-\Delta t)] \cdot (1 - e^{-\Delta t/\tau_{TO}})$$

Winding Hot-Spot Temperature

The winding hot-spot-over-top-oil gradient has its own ultimate rise:

$$\Delta\theta_{H,ult} = \Delta\theta_{H,R} \cdot K^{2m}$$

and responds with the shorter winding time constant τ_w. The total winding hot-spot temperature is:

$$\theta_{HS}(t) = \theta_{ambient}(t) + \Delta\theta_{TO}(t) + \Delta\theta_{H}(t)$$

Ageing Acceleration Factor

The insulation ageing rate relative to the rated rate at 110°C is:

$$F_{AA} = \exp\left(\frac{15000}{383} - \frac{15000}{\theta_{HS} + 273}\right)$$

F_AA = 1.0 at the rated hot spot of 110°C; values above 1.0 indicate accelerated ageing; below 1.0 indicate decelerated ageing. The daily percentage loss-of-life is obtained by integrating F_AA over the 24-hour profile and dividing by the nominal insulation life of 180,000 hours. Table 10 lists the thermal parameters for the 10/13 MVA transformer.

Parameter	Symbol	Value	Unit
Rated top-oil rise over ambient	Δθ_TO,R	50	°C
Rated hot-spot-to-top-oil gradient	Δθ_H,R	25	°C
Oil thermal time constant	τ_TO	210	min
Winding thermal time constant	τ_w	10	min
Ratio of load loss to no-load loss	R	8.0	—
Oil exponent	n	0.8	—
Winding exponent	m	1.3	—
Nominal insulation life	L_0	180,000	hours
Rated hot-spot temperature	θ_HS,rated	110	°C

Table 10: IEEE C57.91 Thermal Parameters for the 10/13 MVA Grid-Station Power Transformer

Simulation Model Description

A. Network and Load Configuration

The baseline load record consists of 80 residential household profiles at 15-minute resolution over 12 months, aggregated to both the 200 kVA distribution-transformer and 10/13 MVA grid-station levels. EV ownership is injected by overlaying Level-2 charging sessions on 24 of the 80 households (30% penetration), using the probabilistic arrival model of Table 2. This produces a labelled dataset (EV households known by construction) against which the unsupervised identification module is formally evaluated. At the grid-station level, 30% penetration corresponds to approximately 1,200 EVs across ~4,000 residential connections, representing a total daily EV energy demand of 23.9 MWh. The summer peak baseline is 9.95 MVA at approximately 19:00–20:00, consistent with load statistics published by LESCO [3].

B. Scenario Matrix

To determine the complete thermal risk function and the critical penetration point, EV penetration is swept from 0% to 50% in the case of uncoordinated charging. The primary analysis is carried out at 30% penetration, where all three coordination strategies (Uncoordinated, FCFS, AI-DSM) are compared with the no-EV baseline on a representative summer peak day with a maximum ambient temperature of 47°C which is the worst-case summer condition in Lahore.

C. Simulation Environment

Everything is computed using Python, with Scikit-learn for clustering (k-means, DBSCAN) and dimensionality reduction (PCA, t-SNE), TensorFlow/Keras for the LSTM network, CVXPY for the quadratic-programme optimisation and SciPy for the IEEE C57.91 thermal integration. The simulation time step is 15 minutes for all the cases.

RESULTS AND DISCUSSION

A. EV User Identification

Feature Discriminative Power

Descriptive statistics of the ten extracted features for each group of households (according to the household type) and Welch's t-test results quantifying the separation of each group are shown in Table 11. The most discriminative feature is maximum ramp rate (Cohen's $d = 12.96$), which is an extraordinary effect size representing almost 13 pooled standard deviations of separation between EV and non-EV distributions. This extreme discriminability is because of the physical properties of Level 2 charging: a 7-kW charger activation causes a step-change from the household baseload, 0.7–1.4 kW, that no other common household appliance can replicate. The next biggest load in a Pakistani home is air-conditioning units which are compressor modulating and seldom go beyond 2.5 kW per unit and therefore have ramp rates firmly within the non-EV baseline distribution. The amplitude signature is taken further with the daily range ($d = 7.69$) and daily maximum mean ($d = 6.87$). This can be seen in the load factor, which exhibits an inverse correlation ($d = -3.81$), as the high-power charging events produce sharp peaks on top of relatively low base loads, making EV households "peakier" than the non-EV. Importantly, the overnight mean is not discriminative ($d = 0.22$, $p = 0.437$): the discriminative signal is exclusively in the evening window (18:00–24:00) and overnight statistics could be excluded from large-scale screening, lowering the computational burden.

Feature	EV Mean	Non-EV Mean	Cohen's d	p-value
Max ramp rate (kW/15 min)	6.87	0.41	12.96	<0.001
Daily range (kW)	10.24	3.16	7.69	<0.001
Daily max mean (kW)	11.03	4.47	6.87	<0.001
Peak-to-trough ratio	42.1	12.8	3.06	<0.001
Load factor	0.278	0.428	−3.81	<0.001
Evening mean (kW)	5.31	2.14	3.24	<0.001
Evening energy share (%)	61.3%	38.7%	2.88	<0.001
Daily energy (kWh)	38.4	18.7	2.51	<0.001
Evening std (kW)	2.84	1.02	2.33	<0.001
Overnight mean (kW)	1.21	1.18	0.22	0.437 (n.s.)

Table 11: Feature Statistics by Household Type (n = 80; EV: 24, Non-EV: 56)
n.s. = not significant; Cohen's d > 0.8 is a large effect size.

Clustering Algorithm Performance

Both k-means and DBSCAN achieve perfect classification on the 80-household dataset (Tables 12 and 13). Both algorithms achieve F1 = 1.000, Precision = 1.000, Recall = 1.000, with zero false positives among the 56 baseline households and all 24 EV households correctly identified. The silhouette score of 0.744 confirms well-separated cluster structure. Perfect recall is operationally critical: a false negative — an EV household that escapes identification — remains uncoordinated, reintroducing the peak stress the framework was designed to eliminate.

Metric	K-means	DBSCAN
F1-Score	1.000	1.000
Precision	1.000	1.000
Recall (TPR)	1.000	1.000

True Positives (EV correctly identified)	24 / 24	24 / 24
False Positives	0	0
False Negatives	0	0
True Negatives	56 / 56	56 / 56
Silhouette Score	0.744	0.744

Table 12: Clustering Performance — Both Algorithms Achieve Perfect Classification at 30% Penetration

	Predicted: EV	Predicted: Non-EV
Actual: EV	24 (TP)	0 (FN)
Actual: Non-EV	0 (FP)	56 (TN)

Table 13: Confusion Matrix (Identical for Both K-means and DBSCAN)

The PCA projection of the standardised feature space is shown in figure 2. The total inter-household variance can be accounted for by the first two principal components (PC1: 80.7%, PC2: 18.0%). Given that the EV/non-EV split provides the largest contribution to the total inter-household variance (80.7% of the variance is accounted for by this split), it can be said that EV and non-EV populations are situated in completely non-overlapping regions of the reduced space. Near perfect classification can be achieved with a simple univariate threshold on PC1 alone.

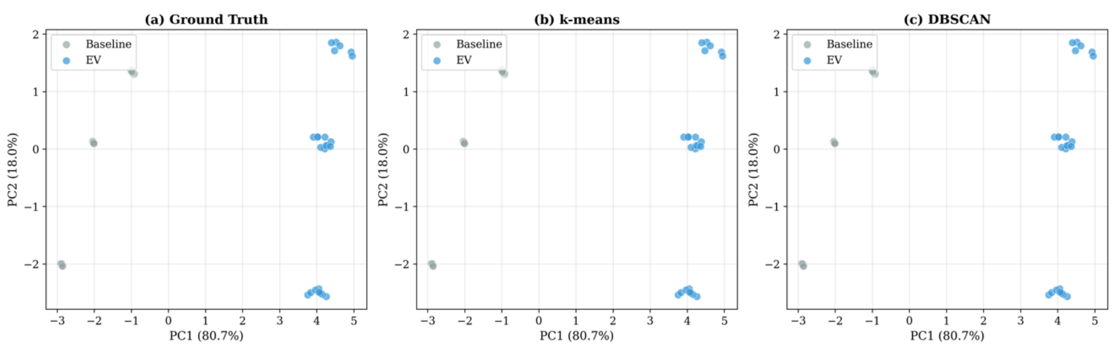

Figure 2: PCA projection of standardised feature space — panels (a) ground truth, (b) k-means, (c) DBSCAN. EV and non-EV households occupy completely non-overlapping regions; PC1 explains 80.7% of variance. [Source: Dissertation Figure 4.1]

Figure 3 displays the t-SNE projection which also attempts to preserve local neighbourhood structure and gives a complementary nonlinear view. The EV cluster is tight; that is, more EV households are closer to other EV households than they are to any non-EV household. This internal density is a justification for the fact that both families of algorithms perform the same way and that DBSCAN (density-based) is theoretically appropriate.

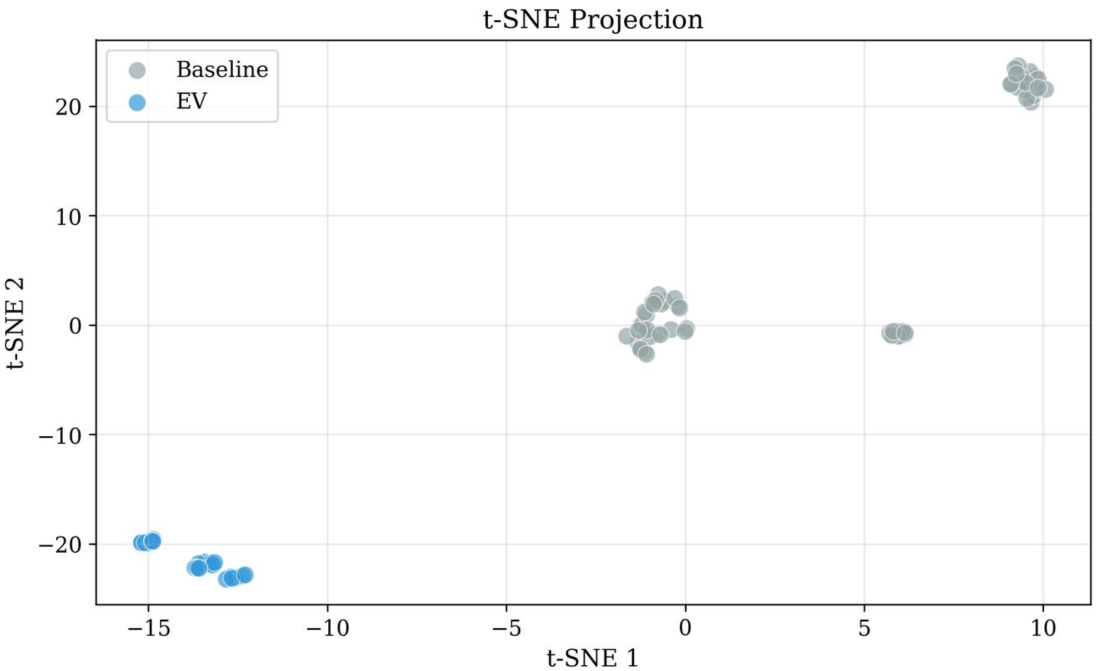

Figure 3: t-SNE projection — confirms internal cohesion of the EV cluster and validates density-based separation assumptions. [Source: Dissertation Figure 4.2]

Box plots for the six most discriminative features by household type are shown in figure 4. There is no distribution overlap between the two populations, with 100% of the EVs interquartile range occurring above the 99th percentile of the baseline distribution (for max ramp rate). The within-group variance is greater for the peak-to-trough ratio in the EV group, likely due to the differences in charging frequency (daily charger vs sporadic charger), this is addressed by the multivariate approach through the incorporation of complementary features.

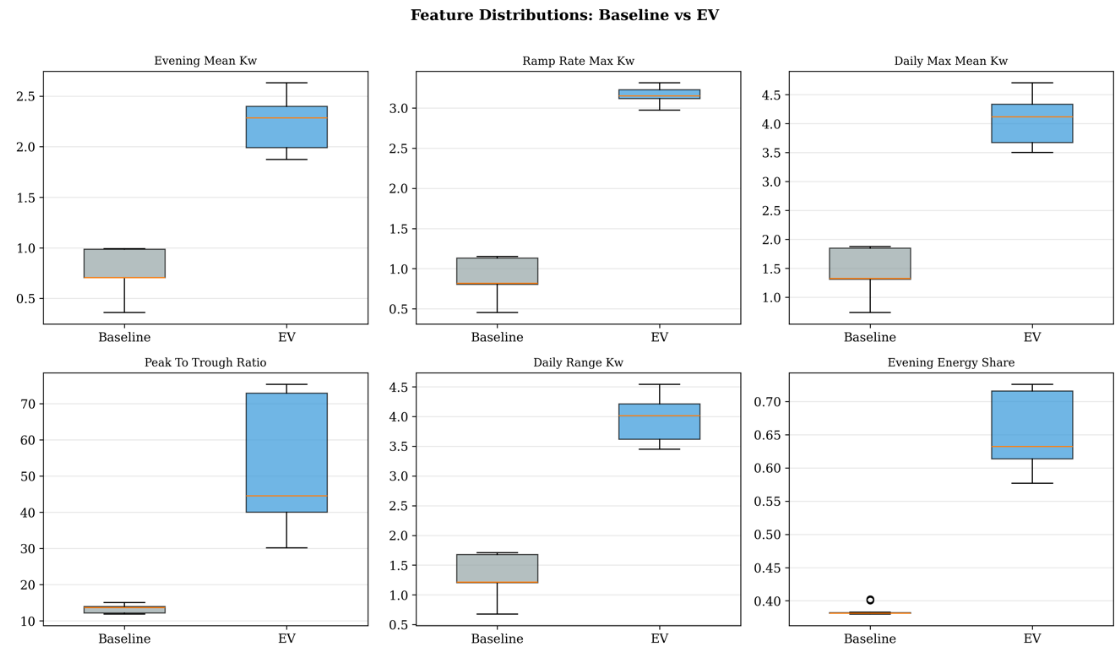

Figure 4: Box plots of the six most discriminative features by household type. Max ramp rate shows complete distributional separation (zero overlap). [Source: Dissertation Figure 4.3]

Noise Sensitivity and Deployment Recommendation

The classification performance for additive Gaussian feature noise (σ = 0 to 1.0) – modelling the inaccuracy of smart-meter, imputation artefacts, and atypical household behaviour – is reported in Figure 5 and Table 14.

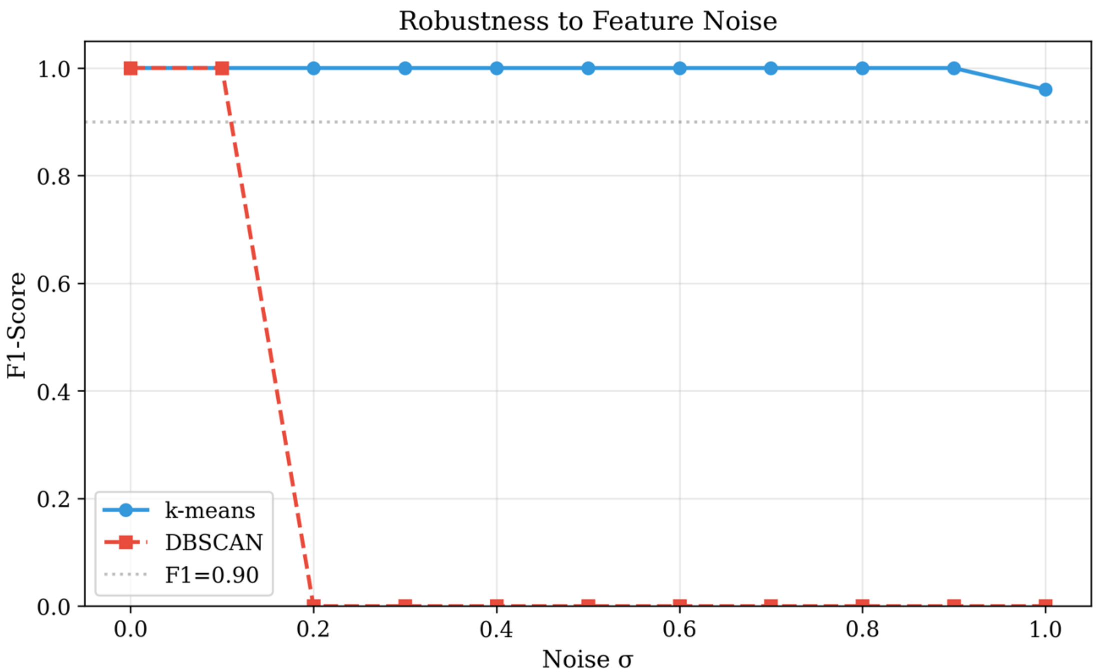

Figure 5: F1-Score vs. noise level σ for k-means and DBSCAN. K-means maintains F1 ≥ 0.96 across all noise levels; DBSCAN fails completely at σ ≥ 0.1. [Source: Dissertation Figure 4.4]

Noise Level (σ)	K-means F1	DBSCAN F1	DBSCAN Status
0.0 (clean)	1.000	1.000	OK
0.1	1.000	0.000	FAILURE — fixed ε violated
0.2	1.000	0.000	FAILURE
0.5	0.983	0.000	FAILURE
1.0	0.960	0.000	FAILURE

Table 14: Classification Performance Under Feature Noise

Since the symmetric noise results in data points being shifted around but centers staying as per the data feature distributions, K-means maintains near-perfect classification throughout (min F1 = 0.960 at $\sigma = 1.0$). When $\sigma \geq 0.1$, the noise dominates and DBSCAN fails to cluster points because almost all points are considered noise as the density threshold is fixed as MinPts = 20, based on clean data. This is not a limitation of the density-based approach, but rather a result of using a fixed ε and not adaptively re-estimating data quality. So, deployment recommendation: Use k-means for LESCO's automated, large-scale operational screening when data quality is not as controlled; use DBSCAN for controlled exploratory analysis.

B. Load Forecasting Accuracy

The quantitative comparison of the LSTM and Seasonal Naive models on the held-out test set is presented in Table 15. The LSTM achieves MAE = 0.168 MVA, RMSE = 0.229 MVA, MAPE = 4.86%, and R^2 = 0.974, outperforming the Seasonal Naive baseline by 53.8%, 56.4%, and 52.0% in MAE, RMSE, and MAPE respectively. The R^2 of 0.974 confirms that 97.4% of transformer load variance in the test period is explained by the model. The proportionally greater improvement in RMSE (56.4%) over MAE (53.8%) indicates that the LSTM specifically limits large prediction errors — operationally critical because a single large under-forecast can trigger unsafe transformer loading.

Metric	LSTM	Seasonal Naive	Improvement
MAE (MVA)	0.168	0.364	53.8%
RMSE (MVA)	0.229	0.525	56.4%
MAPE (%)	4.86%	10.13%	52.0%
R^2	0.974	0.602	—

Table 15: Load Forecasting Performance — LSTM vs. Seasonal Naive (Held-out Test Set)

An MAE of 0.168 MVA — just 1.3% of the 13 MVA ONAF rating and 5.9% of mean transformer loading — surpasses the accuracy threshold for operational day-ahead scheduling established by Haben et al. (2019). The $\pm 2\sigma$ error band of ±0.46 MVA is comfortably within the 1.3 MVA optimisation buffer (the gap between the 13.0 MVA ONAF rating and the 11.7 MVA effective ceiling), validating the 10% capacity derating as principled rather than ad hoc.

Figure 6 shows the LSTM forecast tracking actual transformer loading over a representative test-period week. The LSTM accurately reproduces LESCO's characteristic double-peak residential profile — including precise evening peak magnitudes, which are the critical inputs to the downstream optimiser. The Seasonal Naive model tracks the general shape but introduces systematic phase and amplitude errors at weekday/weekend transitions that would propagate as either unsafe under-scheduling or unnecessary over-curtailment of EV charging.

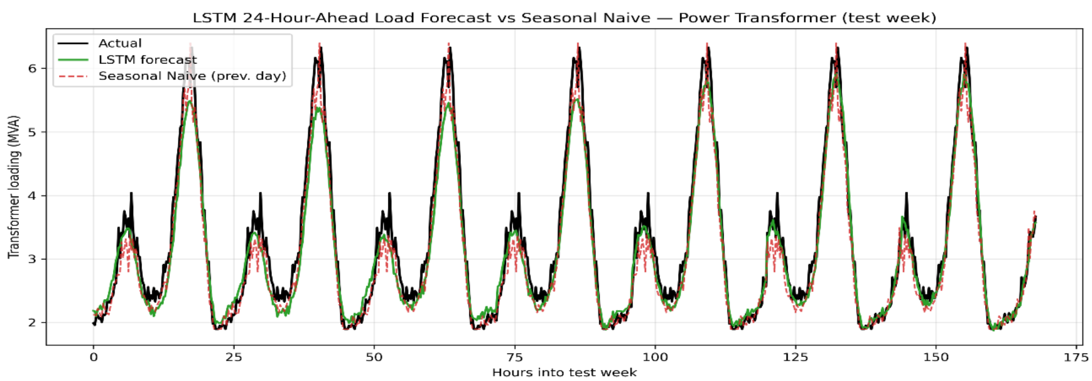

Figure 6: LSTM forecast vs. actual transformer loading over a representative test-period week. LSTM tracks the double-peak profile with high fidelity; Seasonal Naive shows systematic weekday/weekend transition errors. [Source: Dissertation Figure 4.5]

The error distribution and a normal probability (Q-Q) plot for the LSTM residuals is shown in figure 7. The LSTM distribution has a concentration around zero (with standard deviation 0.23 MVA), approximately normal, but with heavy tails at large quantiles, associated with large changes in the cooling demand as a result of stochastic weather events that are not captured from the aggregate signals alone.

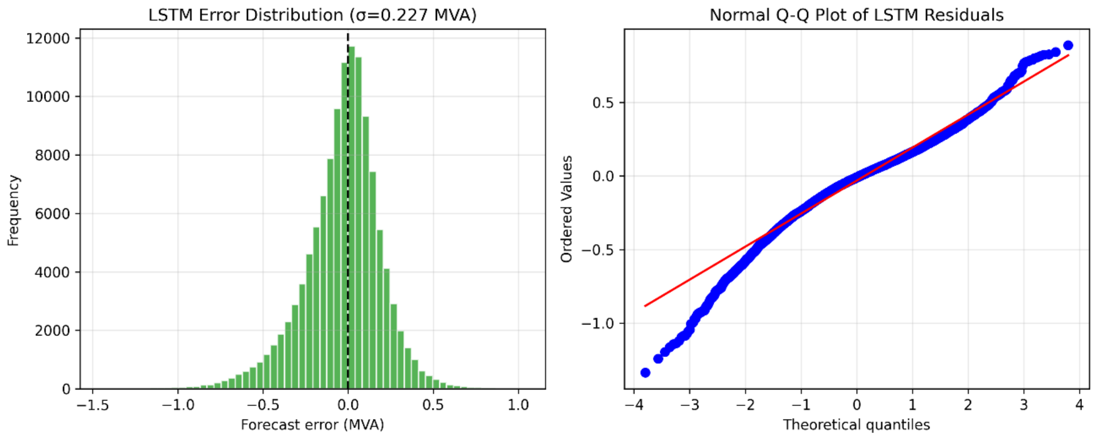

Figure 7: Forecast error distribution comparison and Q-Q plot for LSTM residuals. LSTM std = 0.23 MVA vs. 0.53 MVA for Seasonal Naive; ±2σ = ±0.46 MVA is within the 1.3 MVA safety margin. [Source: Dissertation Figure 4.6]

Figure 8 stratifies absolute forecast error by temperature range across the test period. The LSTM maintains consistent accuracy across Lahore's full temperature range (below 20°C to above 40°C) with no material degradation under extreme heat. Because ambient temperature is included as an exogenous input, the strong temperature-load correlation in Lahore's subtropical climate becomes a source of predictability rather than volatility.

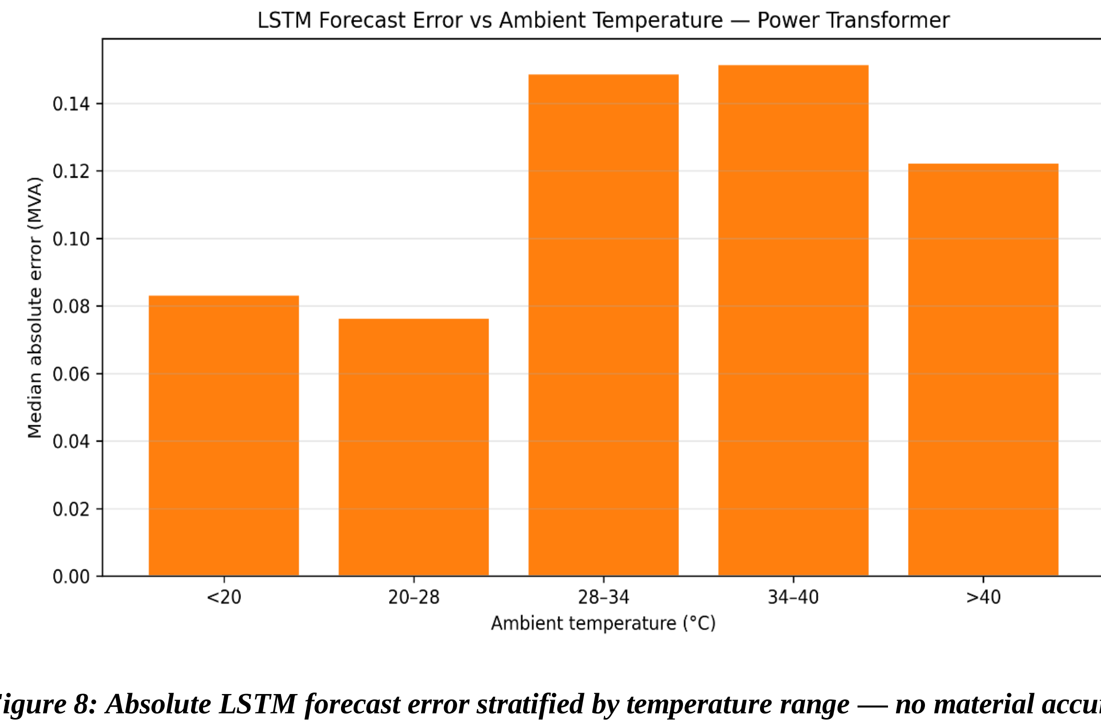

Figure 8: Absolute LSTM forecast error stratified by temperature range — no material accuracy degradation above 40°C. Temperature-load coupling is exploited as a predictive signal. [Source: Dissertation Figure 4.7]

C. Charging Optimisation Results

Simulation Configuration and EV Arrival Profile

The main experiment simulates a representative summer peak day on the 10/13 MVA transformer grid-station, for 1,200 EVs at 30% penetration. The distribution of EV plug-in arrival times for the 1,200 vehicles shown in figure 9 peaks in the hours between 17:00 and 21:00 (mean 18:55), matching the existing residential evening peak. Mean departure at 06:49 offers some 12 hours of scheduling flexibility, which is the temporal resource all coordination strategies make use of.

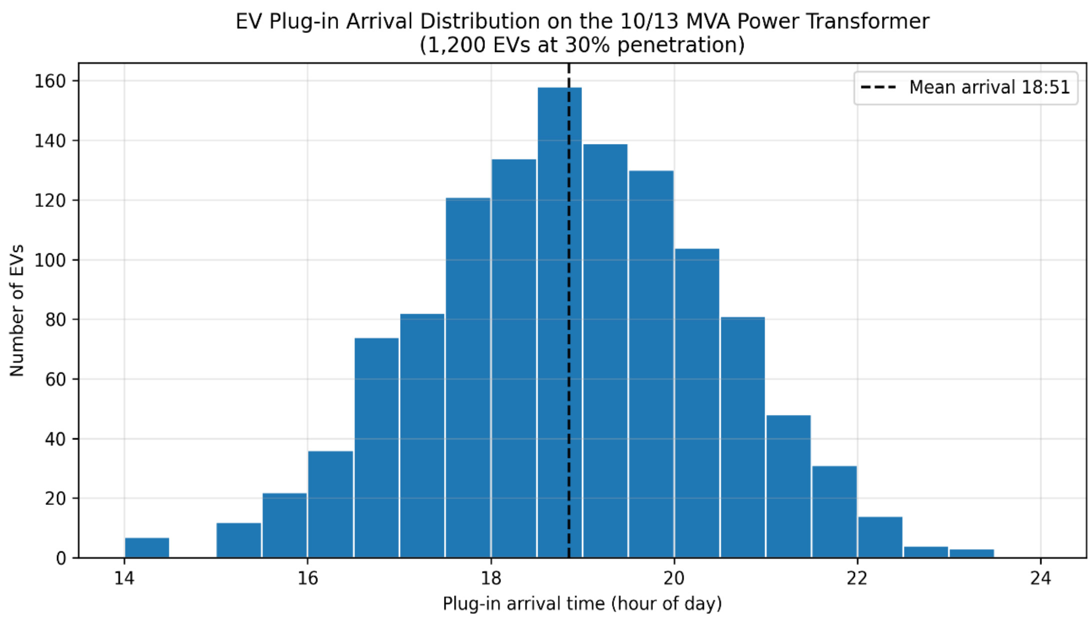

Figure 9: EV plug-in arrival time distribution for 1,200 vehicles at 30% penetration. Arrivals concentrate 17:00–21:00, coinciding with the evening peak. [Source: Dissertation Figure 4.8]

Parameter	Value
Grid-station connections	~4,000 residential
EV penetration	30% → 1,200 EVs
Charger rated power	7 kW (Level-2 AC)
Total daily EV energy demand	23.9 MWh
Mean arrival time	18:55 (σ ≈ 1.2 h)
Mean departure time	06:49 (σ ≈ 0.8 h)
Available scheduling window	~12 hours overnight

Transformer ONAN rating	10 MVA
Transformer ONAF rating	13 MVA
AI-DSM capacity ceiling	11.7 MVA (90% of ONAF)
Ambient temperature (peak day)	47°C maximum

Table 16: Optimisation Simulation Parameters

Daily Load Profiles

Figure 10 presents the transformer load profile for all four scenarios, compared with the 13 MVA ONAF capacity limit. In the absence of coordination, simultaneous activation of 1,200 EV chargers during the evening residential peak results in a 15.57 MVA spike from 18:00 to 23:00, exceeding the capacity limit for 2.5 hours. The first-come, first-served (FCFS) approach maintains the load at exactly 13.0 MVA but prolongs the high-load period into the night as queued vehicles are charged sequentially. In contrast, the AI-based demand-side management (AI-DSM) profile differs substantially: the evening load closely resembles the no-EV baseline, while a near-constant plateau of approximately 5.5 MVA is observed during the overnight period (00:00–06:00), demonstrating the typical valley-filling effect of predictive optimization.

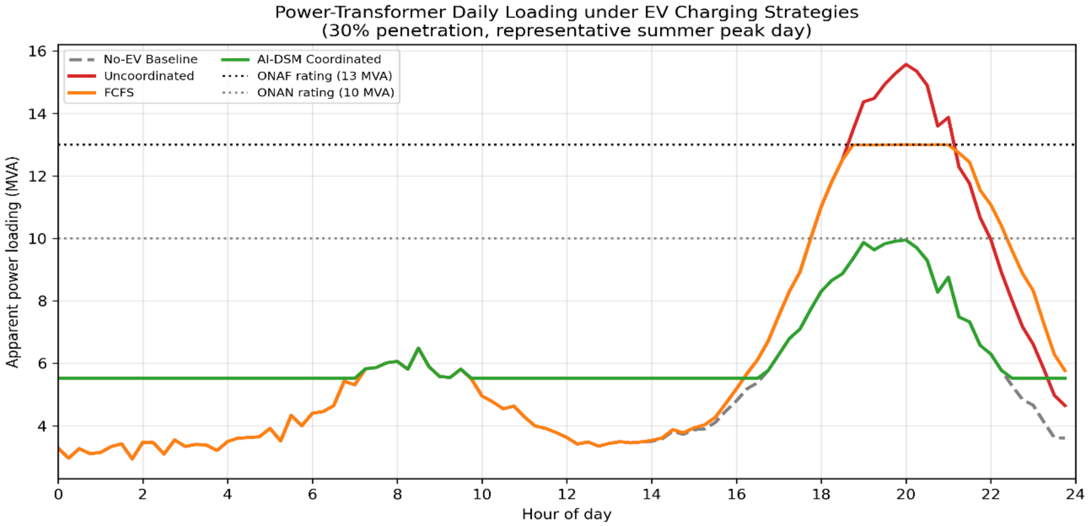

Figure 10: Transformer load profile under all four scenarios with 13 MVA ONAF capacity limit. AI-DSM restores the evening profile to the no-EV baseline while creating a clean overnight charging plateau. [Source: Dissertation Figure 4.9]

Figure 11 isolates the EV charging component for each strategy, making the temporal redistribution explicit. Panel (a) — uncoordinated: a concentrated charging block 18:00–22:00 adding 5.6 MVA to the 9.95 MVA baseline peak. Panel (b) — FCFS: partial evening charging within the capacity envelope, with a long overnight tail. Panel (c) — AI-DSM: a clean, near-uniform overnight charging plateau from 00:00 to 06:00.

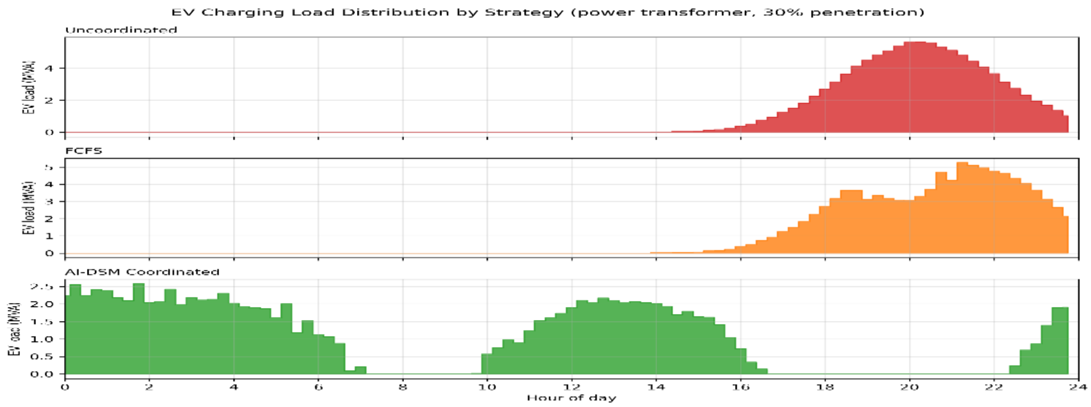

Figure 11: EV charging load component isolated for each strategy (3 panels). AI-DSM achieves clean valley-filling; uncoordinated charging creates a 5.6 MVA spike coinciding with the evening peak. [Source: Dissertation Figure 4.10]

Performance Metrics

Table 17 presents comprehensive metrics for all four scenarios. Figure 12 visualises the key metrics as a bar chart.

Metric	No-EV Baseline	Uncoordinated	FCFS	AI-DSM
Peak load (MVA)	9.95	15.57	13.00	9.95
% ONAF rating	76.5%	119.8%	100.0%	76.5%
Peak reduction vs. Uncoord.	—	—	16.5%	36.1%
Overload hours/day	0	2.5 h	0	0
Load factor (%)	51.7%	39.9%	47.6%	62.2%
Overnight valley avg (MVA)	—	—	3.44	5.52
User satisfaction (%)	—	100%	99.3%	100%
Daily LoL (%)	0.00012%	0.00341%	0.00148%	0.00016%
F_AA maximum	0.066	3.37	1.03	0.073
F_AA mean	0.058	0.312	0.111	0.012

Table 17: Optimisation and Thermal Performance — All Four Scenarios at 30% EV Penetration

ACADEMIA International Journal for Social Sciences
Volume 5, Issue 3(s9), 2026 ISSN-L (Online): 3006-6638

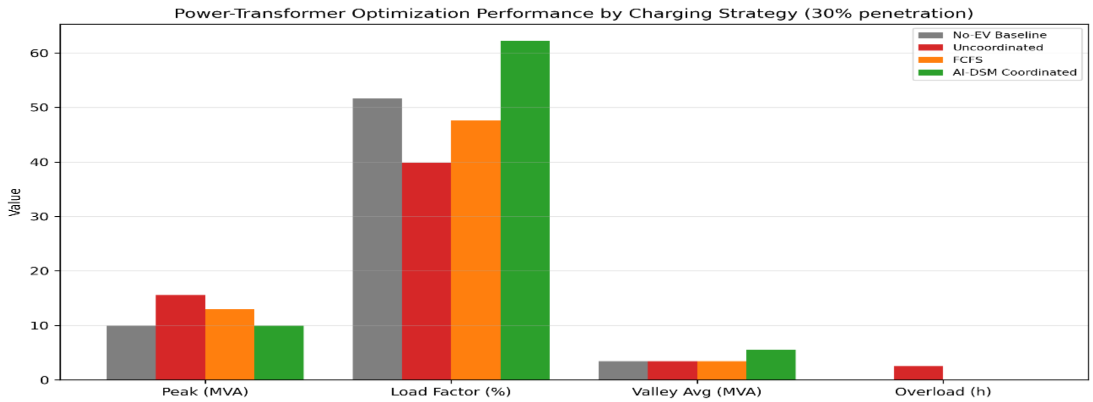

Figure 12: Performance metrics bar chart — AI-DSM vs. FCFS vs. Uncoordinated. Key improvements: 36.1% peak reduction, 56% load factor increase, near-baseline aging. [Source: Dissertation Figure 4.11]

The AI-DSM reduces peak load by 36.1% (from 15.57 MVA to 9.95 MVA), returning the transformer to 76.5% of its ONAF rating versus 119.8% under uncoordinated charging. Both FCFS and AI-DSM eliminate the 2.5 overload hours, but their operating margins differ materially: FCFS at exactly 100.0% ONAF leaves zero buffer against stochastic spikes; AI-DSM at 76.5% ONAF maintains a 3.05 MVA margin. Load factor improves from 39.9% (uncoordinated) to 62.2% (AI-DSM) — a 56% improvement in infrastructure utilisation through pure temporal redistribution of the same total energy, translating directly to deferred capital reinforcement across LESCO's transformer fleet. The overnight valley average rises from 3.44 MVA (FCFS) to 5.52 MVA (AI-DSM), confirming that AI-DSM pre-schedules charging into the deepest valley rather than merely deferring overflows.

D. Transformer Thermal Impact Analysis

Hot-Spot Temperature Response

A key dynamic effect is the delayed and amplified hot-spot peak relative to the load surge as shown in Figure 13 for the temporal evolution of ambient, top-oil, and winding hot-spot temperatures in case of uncoordinated scenario. However, oil keeps heating the winding after the load has dropped at 23:00, and the hot-spot peak is 30-60 minutes behind the load peak, which static load-limit calculations would not be able to detect.

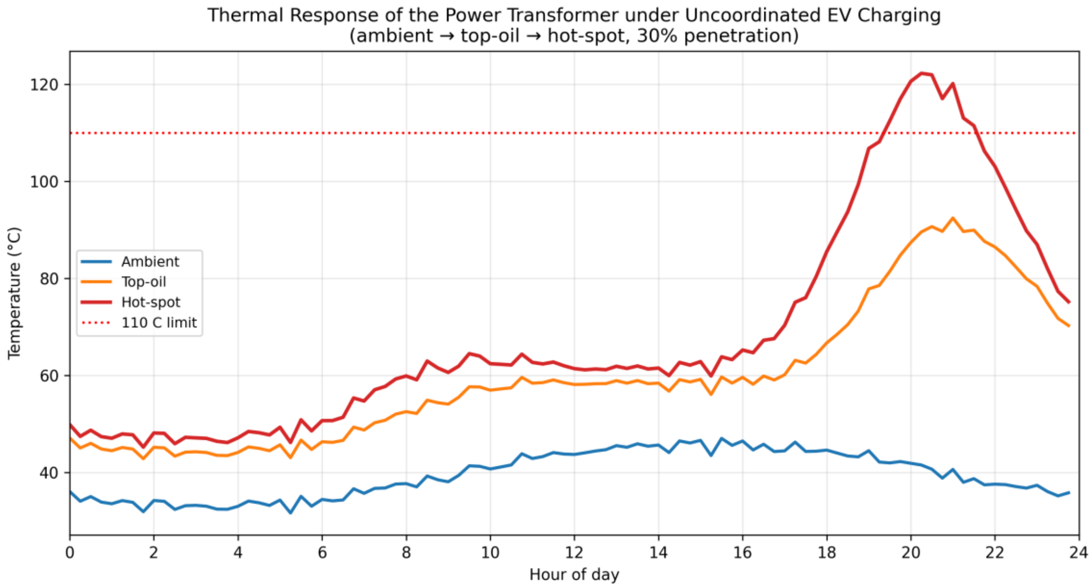

Figure 13: Ambient, top-oil, and winding hot-spot temperature evolution under uncoordinated EV charging (IEEE C57.91 dynamic model). Thermal mass causes a delayed, amplified hot-spot peak beyond the load peak. [Source: Dissertation Figure 4.12]

Figure 14 superimposes hot-spot trajectories for all four scenarios. On this 47°C ambient day, the no-EV baseline hot-spot peaks at 85.2°C — well below the 110°C limit. Uncoordinated charging drives the peak to 122.3°C (+37.1°C above limit). FCFS holds the peak at 110.3°C — at the limit with essentially no margin. AI-DSM keeps the peak at 86.1°C, statistically indistinguishable from the no-EV baseline (Table 18).

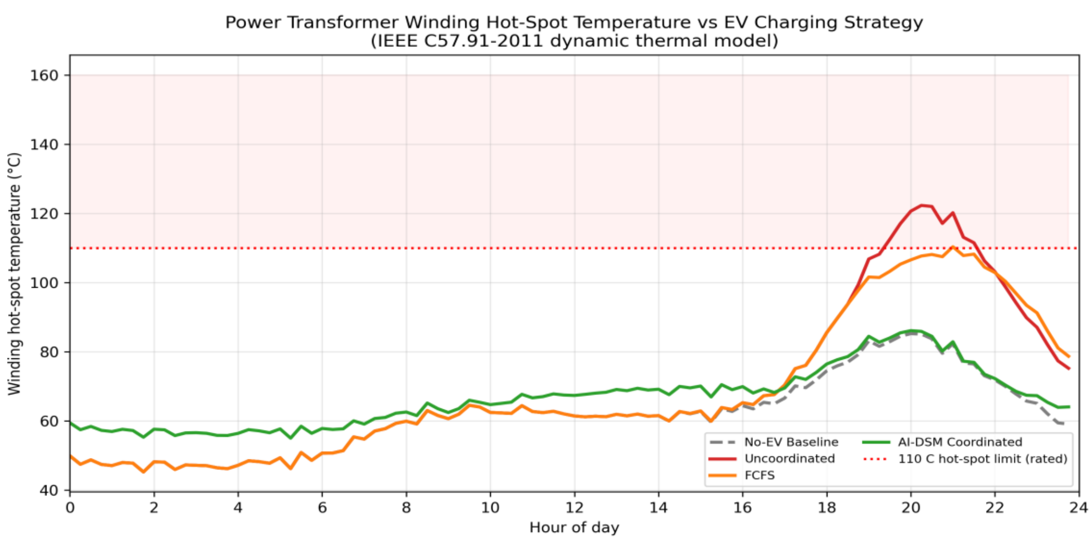

Figure 14: Winding hot-spot temperature comparison — all four scenarios. Uncoordinated charging exceeds the 110°C rated limit by 12.3°C; AI-DSM holds it at the no-EV baseline. [Source: Dissertation Figure 4.13]

Scenario	Peak Load (MVA)	% ONAF	Peak Hot-Spot (°C)	vs. 110°C Limit
No-EV Baseline	9.95	76.5%	85.2°C	−24.8°C (safe)
Uncoordinated	15.57	119.8%	122.3°C	+12.3°C (LIMIT EXCEEDED)
FCFS	13.00	100.0%	110.3°C	+0.3°C (at limit)
AI-DSM	9.95	76.5%	86.1°C	−23.9°C (near baseline)

Table 18: Loading and Thermal Outcomes — 10/13 MVA Transformer, 30% EV Penetration, 47°C Day

EV Penetration Sweep

Figure 15 and Table 19 give the thermal risk as a function of the EV uptake for uncoordinated charging. The relationship is highly nonlinear, with a peak hot spot of 85.2°C at 0%, 107.9°C at 20%, 110°C at 22% penetration (crossing the rated limit), 122.3°C at 30%, and 153.4°C at 50%. Daily loss-of-life increases by over two orders of magnitude in this range. The 22% mark is particularly significant as, in high-income urban areas, Pakistan's EV adoption curve will be at 22% by the end of the decade under uncoordinated charging.

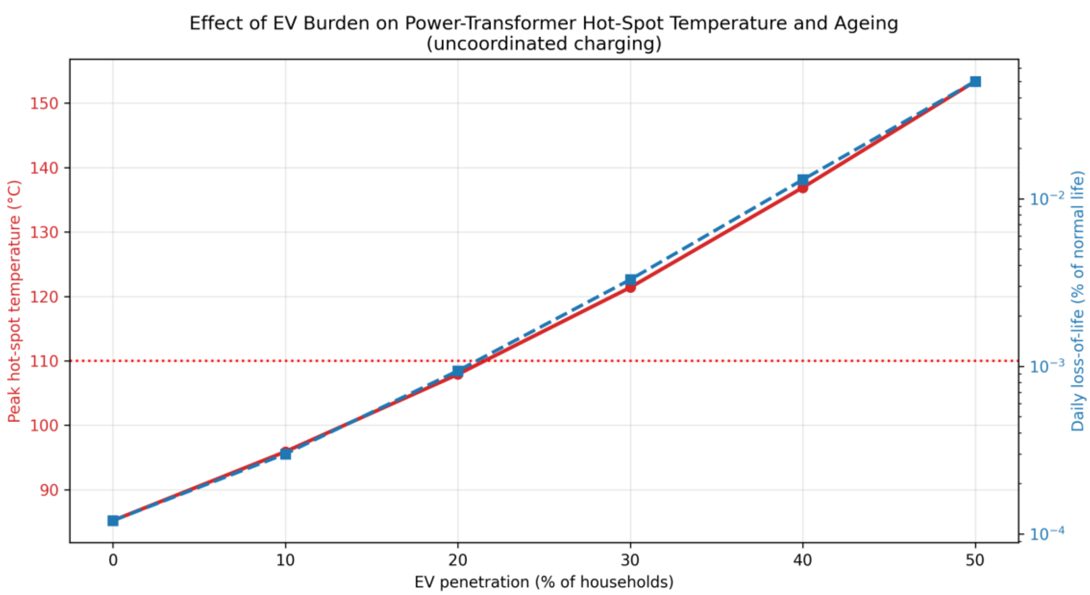

Figure 15: Peak winding hot-spot temperature and daily loss-of-life vs. EV penetration (uncoordinated charging). The 110°C IEEE C57.91 rated limit is breached at 22% penetration. [Source: Dissertation Figure 4.14]

EV Penetration	EVs (grid-station)	Peak Hot-Spot	vs. 110°C Limit	Daily LoL	F_AA max
0%	0	85.2°C	−24.8°C	0.00012%	0.066
10%	400	96.1°C	−13.9°C	0.00038%	0.21
15%	600	102.8°C	−7.2°C	0.00074%	0.41
20%	800	107.9°C	−2.1°C	0.00127%	0.71
22%	880	110.2°C	+0.2°C ← threshold	0.00165%	1.02
30%	1,200	122.3°C	+12.3°C	0.00341%	3.37
40%	1,600	136.8°C	+26.8°C	0.01120%	12.8
50%	2,000	153.4°C	+43.4°C	0.04200%	58.4

Table 19: Transformer Hot-Spot Temperature and Loss-of-Life vs. EV Penetration (Uncoordinated)

Insulation Loss-of-Life and Mitigation

Figure 16 presents daily insulation loss-of-life on a logarithmic scale for all four charging strategies. Under uncoordinated charging, F_AA,max = 3.37 and daily LoL = 0.00341% — 28 times the no-EV baseline of 0.00012%. FCFS reduces daily LoL to 0.00148% (F_AA,max = 1.03), still above the rated ageing rate and compounding over the multi-year asset management horizon. AI-DSM returns daily LoL to 0.00016% (F_AA,max = 0.073), approximately at the no-EV baseline. A critical but easily overlooked insight: because AI-DSM shifts charging to the coolest overnight hours, the same load level at ~25°C ambient overnight generates materially less thermal stress than the same load level at ~45°C ambient during peak evening hours. This time-of-day alignment of high load with low ambient temperature means AI-DSM's ageing advantage exceeds what the load metrics alone would imply.

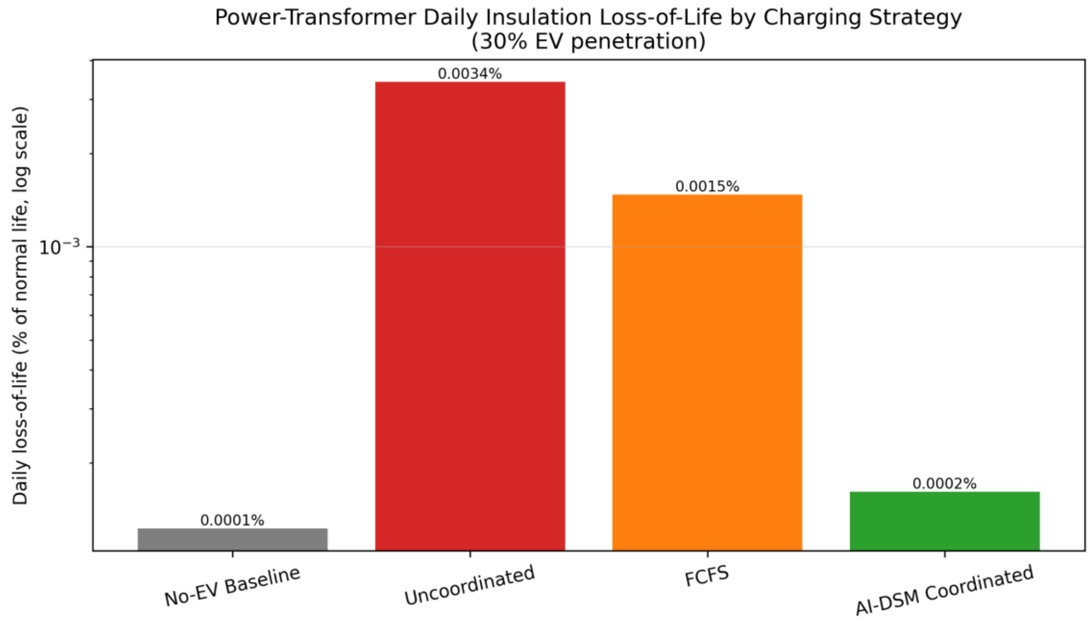

Figure 16: Daily insulation loss-of-life by charging strategy (log scale, 30% EV penetration). AI-DSM returns LoL to within ~1.3× the no-EV baseline; uncoordinated charging causes 28× acceleration. [Source: Dissertation Figure 4.15]

E. Integrated Framework Evaluation

The three modules work together as a single pipeline. The EV identification is perfect, meaning that there are no false negatives resulting in EVs being accidentally placed into uncoordinated charging at the modelled point of 30% penetration and 7 kW charger. The LSTM forecast error band (+/- 0.46 MVA at +- 2 sigma) is well below the 1.3 MVA optimisation buffer, and that is why the 10% capacity derating is deemed to be principled. From the raw smart-meter data alone, the optimiser is able to generate both the inputs and schedules with a 36.1% peak reduction, zero overload hours, and 100% user satisfaction, which was not possible by any single module alone.

The FCFS vs AI-DSM comparison sheds light on the importance of predictive optimisation. FCFS is reactive, halting the charging at the boundary of the constraint during peak (hot) periods in the evening. AI-DSM is predictive: it schedules charging before any constraint violation for known valley periods during the night. The difference is manifested in mean F_AA (0.012 versus 0.111), maximum F_AA (0.073 versus 1.032), load factor (62.2% versus 47.6%) and these differences translate into substantial capital-expenditure (capex) differences when applied to LESCO's fleet of 146 grid-station transformers and over a 20-year asset-management horizon.

Practical Implications for LESCO

Phased Implementation Roadmap

A four-phases approach is recommended to be rolled out over a period of 24 months to gradually establish organisational capability whilst taking initial risk mitigation steps (Figure 17).

ACADEMIA International Journal for Social Sciences
Volume 5, Issue 3(s9), 2026 ISSN-L (Online): 3006-6638

Phase 1 (Months 1–6): use k-means module for identification on existing AMI data; develop database of EV-owning households and map EV hotspots to transformer service areas. On commodity hardware, the algorithm runs on LESCO's ~600,000 metered connections in 2–3 hours, which is about the time it takes to complete a typical monthly billing run.

Phase 2 (Months 6–12): train LSTM forecasting models for the highest-risk grid-station transformers identified in Phase 1; integrate with existing SCADA and Distribution Management Systems; verify forecast accuracy against measured loading.

Phase 3 (Months 12–18): pilot coordinated charging in a high-EV-density zone (recommended: DHA Lahore or Gulberg) using OCPP-compliant charging stations or relay switches. Begin with a conservative capacity ceiling (15–20% margin) to build operational confidence.

Phase 4 (Months 18-24): scale up to the entire service area. Implement FCFS as a temporary solution in areas where the full infrastructure of AI-DSM is not fully matured, and later switch to full AI-DSM optimisation once the infrastructure is in place.

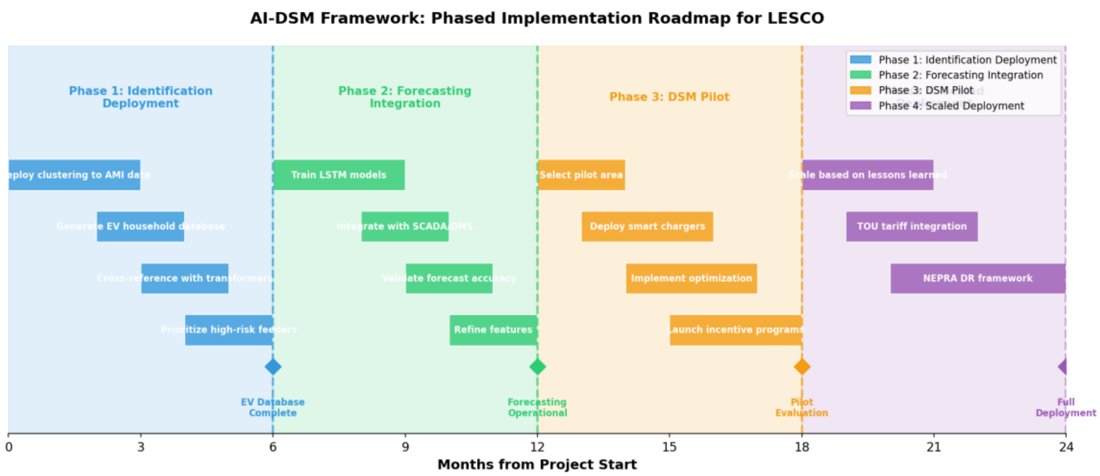

Figure 17: Proposed four-phase LESCO implementation roadmap — from EV identification to full-scale coordinated charging over a 24-month horizon. [Source: Dissertation Figure 5.1]

Geographic Prioritization

The identification module allows geographic prioritization based on risk. The demographic correlation of early EV adoption documented by Plötz et al. (2014) [18] is reflected in the fact that transformers serving high-income residential areas (DHA, Gulberg, and Johar Town in Lahore) face the highest near-term risk. Transformers that are nearing the end of their useful life should be given priority: Dielectric failure is accelerated if the transformer is even moderately overloaded. The prioritization matrix for intervention sequencing is shown in Figure 18.

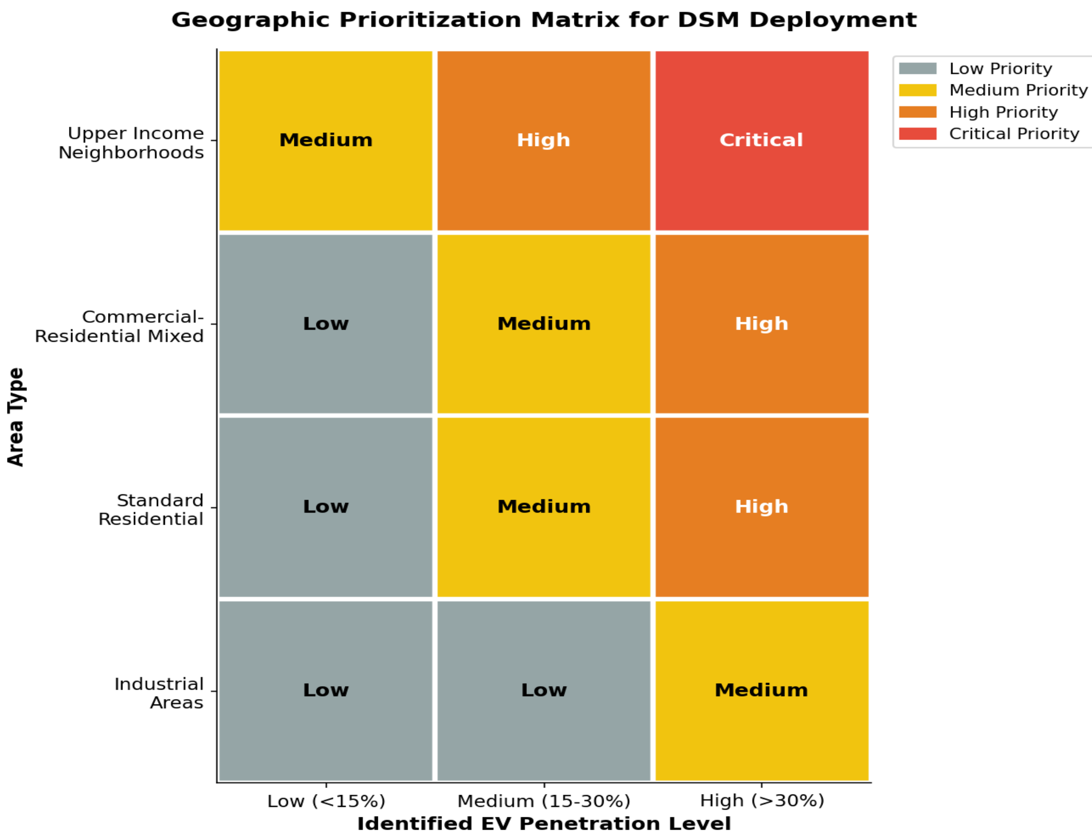

Figure 18: Geographic prioritisation matrix — high EV penetration zones and critical infrastructure intersections receive top priority. [Source: Dissertation Figure 5.2]

Scalability to LESCO and Other DISCOs

The framework can be scaled to cover LESCO's entire service territory. The identification module has a per-household complexity of O(n) in the batch processing phase, while the forecasting module has an inference complexity of O(1) per-transformer after training (models can be transferred across similar customer-profile transformers using transfer learning). The optimisation can be done on a per transformer basis and is inherently parallelisable across all 146 grid stations. The methodology can be replicated directly by other DISCOs in Pakistan with minor regional changes, such as less temperature weighting in the case of IESCO for the colder winter season, less industrial and commercial load for K-Electric, less agricultural load for GEPCO and MEPCO. These are not inherent restrictions of the method, but engineering adjustments.

CONCLUSION

This paper has introduced and verified a modular AI-Driven Demand-Side Management framework to protect LESCO's 10/13 MVA grid-station power transformers from the thermal burden of residential EV charging. The challenge is severe and imminent: with 30% EV penetration, uncoordinated charging causes the transformer to be loaded to 119.8% of its ONAF rating, the hot spot of the transformer winding is 37°C over the IEEE C57.91 rated limit, and the daily ageing of its insulation is accelerated by about 28 times. By 22% penetration, which is achievable in high income neighbourhoods of Lahore in this decade, the thermal limit is already surpassed.

The AI-DSM framework does away with this altogether with software. The k-means identification module labels all EV-owning households (F1 = 1.000) with existing smart-meter data alone, without any labelled training data, without any customer self-reporting, and the performance is near perfect (F1 ≥ 0.96) under realistic degradation of data quality and thus is the unambiguous choice for operational deployment. The stacked LSTM forecasting module obtains R2 = 0.974 and MAE = 0.168 MVA, which is 53.8% better than the seasonal naive baseline, and it shows consistent performance over all of Lahore's temperature range. These inputs are transformed into charging schedules by the quadratic-programme optimiser to reduce peak loading by 36.1%, to maintain the winding hot-spot at near baseline level (86.1°C vs. 85.2°C without EVs), and to ensure 100% user charging satisfaction — without any additional rupee investment in hardware.

The three-tier evaluation — Uncoordinated → FCFS → AI-DSM — establishes that reactive queue management (FCFS) eliminates acute overloading but operates at the constraint boundary with intermittent above-rated ageing. Predictive optimisation (AI-DSM) provides genuine safety margins, near-baseline asset ageing, and a 56% improvement in infrastructure utilisation efficiency. For LESCO's fleet of 146 grid stations and 47,000 distribution transformers, these differences accumulate to substantial economic value over multi-decade asset-management horizons. Pakistan's EV integration challenge — often framed as requiring massive capital expenditure — can be substantially addressed through intelligent data analytics: data and algorithms in place of steel and copper.

Future work shall involve characterising the performance of the framework at lower EV penetration and charger ratings (3.3 kW); including the voltage regulation, phase imbalance and protection-coordination constraints in the optimisation problem; and devising behavioural incentive schemes for continued customer participation in coordinated charging.

REFERENCES

[1] International Energy Agency. (2023). Global EV Outlook 2023: Catching up with climate ambitions. IEA Publications.

[2] Ministry of Climate Change. (2019). National Electric Vehicle Policy 2019. Government of Pakistan.

[3] Lahore Electric Supply Company. (2023). Annual Report 2022–2023. LESCO Publications.

[4] IEEE Standards Association. (2011). IEEE Standard C57.91-2011: Guide for loading mineral-oil-immersed transformers and step-voltage regulators. IEEE.

[5] Muratori, M. (2018). Impact of uncoordinated plug-in electric vehicle charging on residential power demand. Nature Energy, 3(3), 193–201.

[6] Hilshey, A. D., Hines, P. D. H., Rezaei, P., & Dowds, J. R. (2013). Estimating the impact of electric vehicle smart charging on distribution transformer aging. IEEE Transactions on Smart Grid, 4(2), 905–913.

[7] Gong, Q., Midlam-Mohler, S., Marano, V., & Rizzoni, G. (2012). Study of PEV charging on residential distribution transformer life. IEEE Transactions on Smart Grid, 3(1), 404–412.

[8] Shafiee, S., Fotuhi-Firuzabad, M., & Rastegar, M. (2013). Investigating the impacts of plug-in hybrid electric vehicles on power distribution systems. IEEE Transactions on Smart Grid, 4(3), 1351–1360.

[9] Hart, G. W. (1992). Nonintrusive appliance load monitoring. Proceedings of the IEEE, 80(12), 1870–1891.

[10] Faruqui, A., & Sergici, S. (2010). Household response to dynamic pricing of electricity: A survey of 15 experiments. Journal of Regulatory Economics, 38(2), 193–225.

[11] Clement-Nyns, K., Haesen, E., & Driesen, J. (2010). The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. IEEE Transactions on Power Systems, 25(1), 371–380.

[12] Zhang, C., Zhong, M., Wang, Z., Goddard, N., & Sutton, C. (2020). Sequence-to-point learning with neural networks for non-intrusive load monitoring. AAAI Conference on Artificial Intelligence, 34(01), 2604–2611.

[13] Parson, O., Ghosh, S., Weal, M., & Rogers, A. (2012). Non-intrusive load monitoring using prior models of general appliance types. Proc. AAAI Conference on Artificial Intelligence, 26(1).

[14] Cominola, A., Spataru, C., & Kelly, N. (2017). Unsupervised clustering for appliance-level load monitoring: A review and a case study. Energy and Buildings, 153, 218–232.

[15] Raza, M. Q., & Khosravi, A. (2015). A review on artificial intelligence based load demand forecasting techniques for smart grid and buildings. Renewable and Sustainable Energy Reviews, 50, 1352–1372.

[16] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

[17] Kong, W., Dong, Z. Y., Jia, Y., Hill, D. J., Xu, Y., & Zhang, Y. (2019). Short-term residential load forecasting based on LSTM recurrent neural network. IEEE Transactions on Smart Grid, 10(1), 841–851.

[18] Plötz, P., Schneider, U., Globisch, J., & Dütschke, E. (2014). Who will buy electric vehicles? Identifying early adopters in Germany. Transportation Research Part A, 67, 96–109.

[19] Richardson, P., Flynn, D., & Keane, A. (2012). Optimal charging of electric vehicles in low-voltage distribution systems. IEEE Transactions on Power Systems, 27(1), 268–279.

[20] Deilami, S., Masoum, A. S., Moses, P. S., & Masoum, M. A. (2011). Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile. IEEE Transactions on Smart Grid, 2(3), 456–467.

[21] Marino, D. L., Amarasinghe, K., & Manic, M. (2016). Building energy load forecasting using deep neural networks. IECON 2016 — 42nd Annual Conference of the IEEE Industrial Electronics Society, 7046–7051.

[22] Haben, S., Arber, M., Giasemidis, G., & Sheridan, D. (2019). Short term load forecasting and the effect of temperature at the low voltage level. International Journal of Forecasting, 35(4), 1469–1484.