

Demand-Side Management of EV Integration with LESCO's Grid Stations



SUPERIOR UNIVERSITY

The thesis submitted in partial fulfilment of the requirements for the degree of

**MASTER OF PHILOSOPHY
IN
ELECTRICAL ENGINEERING**

Submitted by: Mudassar Rafique

Registration No: SU92-MSEEW-S24-004

2024-2026

Supervised by

Dr Manzoor Ellahi

DEPARTMENT OF ELECTRICAL ENGINEERING

THE SUPERIOR UNIVERSITY, LAHORE

DECLARATION

I, Mudassar Rafique, a student of “**MS Electrical Engineering**”, at “**Faculty of Engineering and Technology**”, the Superior University, Lahore, hereby declare that this thesis titled “**Demand-Side Management of EV Integration with LESCO’s Grid Stations**” is my own research work and has not been submitted, published, or printed elsewhere in Pakistan or abroad. Additionally, I will not use this thesis to obtain any degree other than the one stated above.

I fully understand that if my statement is found to be incorrect at any stage, including after the award of the degree, the university has the right to revoke my MS/M.Phil degree.

Signature of Student: M. Mudassar

Name of Student: Mudassar Rafique

Roll Number: SU92-MSEEW-S24-004

Date: 06 July 2026

PLAGIARISM UNDERTAKING

I solemnly declare that the research work presented in this thesis, titled “**Demand-Side Management of EV Integration with LESCO’s Grid Stations**”, is solely my research work, and that the entire thesis has been completed by me, with no significant contribution from any other person or institution. Any small contribution, wherever taken, has been duly acknowledged.

I understand the zero-tolerance policy of the HEC and Superior University, Lahore, towards plagiarism. Therefore, I, as the author of the above-titled thesis, declared that no portion of my thesis has been plagiarized and that all material used from other sources has been properly acknowledged, cited, and referenced.

I undertake that if I am found guilty of any formal plagiarism in the above-titled thesis, even after the award of MS/M.Phil degree, the University reserves the right to revoke my degree, and HEC and the University have the right to publish my name on the HEC/University website for submitting a plagiarized thesis.

Signature of Student: M. Qudusar

Signature of Supervisor: Manzoor Ellahi

Supervisor Name: Dr Manzoor Ellahi

CERTIFICATE OF RESEARCH COMPLETION

It is certified that this thesis titled “**Demand-Side Management of EV Integration with LESCO’s Grid Stations**”, submitted by_Mudassar Rafique, Roll No. **SU92-MSEEW-S24-004**, for MS **Electrical Engineering** at “**Faculty of Engineering and Technology**”, Superior University, is an original research work and contains satisfactory material to be eligible for evaluation by the Examiner(s) for the award of the above-stated degree.

No part of this dissertation has been submitted anywhere else for any other degree.

Supervisor’s Signature: _____

Supervisor’s Name: Dr. Manzoor Ellahi

Designation: Associate Professor

CERTIFICATE OF APPROVAL

It is certified that the research work contained in this thesis, titled “**Demand-Side Management of EV Integration with LESCO’s Grid Stations**”, was conducted by “**Mudassar Rafique**” under the supervision of “**Dr Manzoor Ellahi**” and is adequate for the award of “**Master of Philosophy in Electrical Engineering**”.

No part of this thesis has been submitted elsewhere for any other degree. This thesis is submitted to the “**Faculty of Engineering and Technology**”, The Superior University, Lahore, in partial fulfilment of the requirements for the degree.

Student Name: Mudassar Rafique

Signature: M. Mudassar

Examination Committee:

a) **Session Chair:** Dr. Muhammad Zafrullah

Signature: _____

b) **External Examiner**

Name:

Signature: _____

Designation:

University Name:

c) **Internal Examiner**

Name: Dr. Irfan Mushtaq

Signature: _____

Designation: Assistant Professor

University Name: The Superior University, Lahore

d) **Supervisor Name:** Dr. Manzoor Ellahi

Signature: _____

e) **Head of Department:**

Signature: _____

f) **Name of the Dean:** Dr. Muhammad Zafrullah

Signature: _____

g) **COE:** Dr. Muhmmad Haris

Signature: _____

DEDICATION

To my father, whose unwavering love, patience, and constant encouragement have been the cornerstone of my academic journey. Their sacrifices, understanding, and belief in my abilities provided me with the strength, motivation, and resilience required to overcome challenges and remain committed to my research goals. Throughout the demanding phases of this work, their emotional support and confidence in my aspirations served as a continuous source of inspiration. I also dedicate this thesis to my teachers and mentors, whose guidance, wisdom, and dedication to academic excellence profoundly influenced my intellectual development. Their encouragement to think critically, pursue innovation, and uphold ethical research practices shaped the foundations of my academic and professional growth. The knowledge and values imparted by them have left a lasting impact on my approach to learning and research. Finally, this work is dedicated to all individuals who believe in the transformative power of education and research in advancing society. It stands as a reflection of the collective support, inspiration, and shared commitment that made this academic achievement possible.

ACKNOWLEDGEMENTS

I want to express my deepest and most sincere gratitude to my supervisor, Dr Manzoor Ellahi, for his continuous guidance, insightful supervision, and constructive feedback throughout the entire course of this research. His profound technical knowledge, research experience, and critical perspective played a central role in shaping the conceptual framework, methodological design, and analytical depth of this thesis. His encouragement, availability for discussion, and consistent academic support were instrumental in overcoming research challenges and ensuring the successful and timely completion of this work. I am also sincerely thankful to the faculty members of the Department of Electrical Engineering for providing a strong academic foundation and fostering a supportive and intellectually stimulating research environment. Their lectures, scholarly discussions, and academic mentorship significantly enhanced my understanding of power systems, demand-side management, and smart grid technologies. The knowledge and skills gained through their instruction and guidance have been invaluable to both the theoretical and practical aspects of this research.

Furthermore, I gratefully acknowledge the support and cooperation of professionals and technical staff associated with the Lahore Electric Supply Company (LESCO). Their practical insights into distribution network operations, transformer loading issues, and real-world grid constraints greatly contributed to grounding this research in realistic utility conditions. Their input helped bridge the gap between theoretical modelling and practical implementation, thereby enhancing the applicability and relevance of the proposed framework. I would also like to extend my appreciation to my colleagues and peers for their constructive discussions, collaboration, and moral support throughout the research process. The exchange of ideas, technical discussions, and shared academic experiences was highly beneficial in refining the research methodology, validating assumptions, and improving the overall quality of the analysis presented in this thesis. Finally, I express my heartfelt gratitude to my family for their unwavering support, patience, and encouragement during my academic journey. Their understanding, motivation, and emotional support provided strength, stability, and

confidence is required to remain focused and committed throughout the demanding phases of this research. This achievement would not have been possible without their constant encouragement and belief in my abilities.

Thank you,

Mudassar Rafique

TABLE OF CONTENTS

DECLARATION.....	II
PLAGIARISM UNDERTAKING.....	III
CERTIFICATE OF RESEARCH COMPLETION.....	IV
CERTIFICATE OF APPROVAL	V
DEDICATION	VI
ACKNOWLEDGEMENTS	VII
TABLE OF CONTENTS	IX
LIST OF FIGURES	XIII
LIST OF TABLES.....	XIV
LIST OF ABBREVIATIONS AND ACRONYMS	XVI
ABSTRACT	1
Chapter 1: Introduction.....	3
1.1 Background	3
1.2 Problem Statement	4
1.3 Research Aim	5
1.4 Research Objectives	6
1.5 Research Questions	6
1.6 Scope of the Study.....	7
1.7 Significance of Study	7
1.8 Definition of Key Terms	8
1.9 Thesis Outline.....	9
Chapter 2: Literature Review	10
2.1 The Impact of EVs on Distribution Networks.....	10
2.2 EV Load Identification Techniques (NILM).....	11

2.3 Short-Term Load Forecasting (STLF).....	12
2.4 Summary and Research Gap	13
Chapter 3: Methodology	14
3.1 Research Design and Framework Overview.....	14
3.1.2 Research Philosophy and Approach	14
3.2 Data Collection and Preprocessing.....	16
3.2.1 Data Preprocessing Pipeline	18
3.2.2 Data Partitioning Strategy	18
3.3 Module 1: EV User Identification.....	19
3.3.1 k-means Implementation	20
3.3.2 DBSCAN Implementation.....	20
3.3.3 Validation Metrics	21
3.4 Module 2: Load Forecasting.....	21
3.4.1 Model Architecture.....	21
3.4.2 Input Feature Vector.....	22
3.4.3 Training Protocol.....	22
3.4.4 Baseline Model.....	23
3.5 Module 3: Charging Optimization.....	23
3.5.1 Charging Scenarios.....	23
3.5.2 Problem Formulation (AI-DSM).....	24
3.5.3 Constraints:.....	24
3.5.4 Transformer Ageing Assessment.....	25
3.6 Power-Transformer Thermal Modelling (IEEE C57.91-2011).....	25
Chapter 4: Results and Analysis.....	28
4.1 EV User Identification Performance.....	28
4.1.1 Dataset and Feature Extraction	28
4.1.2 Clustering Algorithm Performance	31

4.1.3 Dimensionality Reduction Visualization.....	32
4.1.4 Sensitivity Analysis.....	35
4.1.5 Summary of EV Identification Results.....	36
4.2 Load Forecasting Accuracy	37
4.2.1 Experimental Configuration.....	37
4.2.2 Model Performance Comparison.....	37
4.2.3 Forecast Tracking Analysis	39
4.2.4 Error Distribution Analysis	40
4.2.5 Temperature Sensitivity.....	40
4.2.6 Summary of Forecasting Results.....	41
4.3 Charging Optimization on the 10/13 MVA Power Transformer.....	41
4.3.1 Simulation Configuration	42
4.3.2 Scenario Definitions	43
4.3.3 Daily Load Profile Comparison.....	44
4.3.4 EV Charging Load Distribution	45
4.3.5 Performance Metrics	46
4.3.6 Transformer Loss-of-Life under the Dynamic Thermal Model	48
4.3.7 Comparative Assessment of FCFS versus AI-DSM.....	49
4.3.8 Summary of Optimization Results	50
4.4 Power Transformer (10/13 MVA) Thermal Impact Analysis.....	50
4.4.1 From Loading to Winding Temperature	50
4.4.2 Winding Hot-Spot Temperature Response.....	51
4.4.3 Effect of EV Penetration on Temperature and Ageing.....	52
4.4.4 Loss-of-Life and Mitigation	53
4.4.5 Summary of Power-Transformer Thermal Findings	54
4.5 Integrated Framework Evaluation.....	55
Chapter 5: Discussion	56

5.1 Interpretation of Findings	56
5.1.2 EV Identification Through Unsupervised Learning	56
5.1.3 LSTM-Based Load Forecasting Performance	57
5.1.4 Grid Impact of Coordinated Charging.....	58
5.2 Addressing the "Blindness" Problem.....	60
5.3 Implications for LESCO Implementation	60
5.3.1 Phased Rollout Strategy.....	61
5.3.2 Geographic Prioritization	62
5.4 Scalability Assessment	63
5.4.1 Computational Scalability.....	63
5.5 Limitations.....	64
5.6 Comparison with Prior Work	65
5.7 Summary	66
References	68

LIST OF FIGURES

Figure 3. 1: AI-DSM Framework Conceptual Diagram.....	15
Figure 4. 1: presents the PCA projection in three panels: (a) ground truth labels, (b) k-means predictions, and (c) DBSCAN predictions.	32
Figure 4. 2: presents the t-SNE projection, which preserves local neighborhood structure and provides a complementary nonlinear view.....	33
Figure 4. 3: presents box plots comparing the distribution of the six most discriminative features between the two household types.	34
Figure 4. 4: plots F1-Score against noise level for both algorithms.....	35
Figure 4. 5: Illustrates the LSTM's ability to track actual load over a representative week in the test period, compared with the Seasonal Naive prediction.	39
Figure 4. 6: presents the error distribution comparison and a normal probability (Q-Q) plot for the LSTM residuals.	40
Figure 4. 7: stratifies the absolute forecast error by temperature range across the test period....	41
Figure 4. 8: presents these metrics in bar chart form for visual comparison.....	47
Figure 4. 9: Ambient, top-oil and winding hot-spot temperature evolution under uncoordinated EV charging (IEEE C57.91 dynamic model).	51
Figure 4. 10: Winding hot-spot temperature of the power transformer under each EV charging strategy. Uncoordinated charging exceeds the 110 C rated limit; AI-DSM holds it at the no-EV baseline.	52
Figure 4. 11 Peak winding hot-spot temperature and daily loss-of-life of the power transformer versus EV penetration (uncoordinated charging).:.....	53
Figure 4. 12: Daily insulation loss-of-life of the 10/13 MVA power transformer by charging strategy (log scale, 30% EV penetration).	54
Figure 5. 1: Proposed four-phase implementation roadmap for LESCO, transitioning from identification to full-scale deployment over a 24-month horizon.	61
Figure 5. 2: Geographic prioritization matrix guiding LESCO's DSM deployment. High EV penetration areas and critical infrastructure zones receive top priority.....	63

LIST OF TABLES

Table 3. 1: Network Hierarchy and Scope	16
Table 3. 2: EV Charging Session Parameters	17
Table 3. 3: Dataset Summary	18
Table 3. 4: Data Partitioning	19
Table 3. 5: Features Extracted for Clustering	20
Table 3. 6: Clustering Evaluation Metrics	21
Table 3. 7: LSTM Model Hyperparameters	22
Table 3. 8: Forecasting Evaluation Metrics	23
Table 3. 9: Charging Scenario Definitions	23
Table 3. 10: Optimization Problem Summary	25
Table 3. 11: IEEE C57.91 thermal parameters for the 10/13 MVA power transformer	26
Table 4. 1: Dataset Summary for Clustering Analysis	29
Table 4. 2: Feature Statistics by Household Type	30
Table 4. 3: Clustering Algorithm Performance Comparison	31
Table 4. 4: Confusion Matrix (Both Algorithms)	31
Table 4. 5: Classification Performance Under Feature Noise	35
Table 4. 6: Data Partition for Forecasting	37
Table 4. 7: Forecasting Performance Metrics	38
Table 4. 8: Optimization Simulation Parameters	42
Table 4. 9: presents the distribution of EV plug-in arrival times across 1,200 vehicles.	43
Table 4. 10: presents the transformer load profile under all three charging scenarios alongside the no-EV baseline and the 13 MVA ONAF capacity limit.....	44
Table 4. 11: isolates the EV charging component to illustrate how each strategy distributes the 1,200-vehicle charging load across 24 hours.	45
Table 4. 12: Optimization Performance Comparison.....	46
Table 4. 13: Loading and thermal outcomes on the 10/13 MVA power transformer (30% EV	

penetration) for all four scenarios.	52
Table 4. 14: Power-transformer loading, hot-spot temperature and loss-of-life versus EV penetration (uncoordinated charging).	53

LIST OF ABBREVIATIONS AND ACRONYMS

AC	Air Conditioning / Alternating Current
AI	Artificial Intelligence
AI	DSM – Artificial Intelligence-Driven Demand-Side Management
AMI	Advanced Metering Infrastructure
ARIMA	Auto-Regressive Integrated Moving Average
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DHA	Defense Housing Authority
DISCO	Distribution Company
DSM	Demand-Side Management
EV	Electric Vehicle
FCFS	First-Come, First-Served
FFNN	Feed-Forward Neural Network
GEPCO	Gujranwala Electric Power Company
HEC	Higher Education Commission
HMM	Hidden Markov Model
HVAC	Heating, Ventilation, and Air Conditioning
IEA	International Energy Agency
IEEE	Institute of Electrical and Electronics Engineers
IESCO	Islamabad Electric Supply Company
kV	Kilovolt
kVA	Kilovolt-Ampere
kW	Kilowatt
kWh	Kilowatt-hour
LESCO	Lahore Electric Supply Company
LoL	Loss of Life
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error

MEPCO	Multan Electric Power Company
MSE	Mean Squared Error
MVA	Mega Volt-Ampere
NEVP	National Electric Vehicle Policy
NILM	Non-Intrusive Load Monitoring
ONAF	Oil Natural Air Forced
ONAN	Oil Natural Air Natural
OSQP	Operator Splitting Quadratic Program (solver)
PCA	Principal Component Analysis
PHEV	Plug-in Hybrid Electric Vehicle
PKR	Pakistani Rupee
PV	Photovoltaic
QP	Quadratic Program
Q-Q (plot)	Quantile-Quantile (plot)
R²	Coefficient of Determination (R-squared)
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network
SARIMA	Seasonal Auto-Regressive Integrated Moving Average
SCADA	Supervisory Control and Data Acquisition
SDT	Secondary Distribution Transformer
SoC	State of Charge
STLF	Short-Term Load Forecasting
ToU	Time of Use
t-SNE	t-distributed Stochastic Neighbor Embedding

ABSTRACT

The rapid penetration of EVs into Pakistan's transportation sector is both an opportunity for sustainable mobility and a major threat to power distribution channels. An increasing strain on the grid of the Lahore Electric Supply Company (LESCO) is caused by uncontrollable and simultaneous processes of EV charging, and they often occur during evening peak demand periods. There is overloading on the transformers due to such uncontrolled operations, instability of the voltage and inefficient usage of energy. Absence of real-time observation, load-sensitive coordination and load scheduling also play a part in these challenges that threaten the resilience and reliability of the existing infrastructure. This study suggests a smart Demand-Side Management (DSM) model based on Artificial Intelligence to organize EV charging in an intelligent way throughout the distribution system of LESCO. It combines machine learning and predictive analytics to facilitate real-time identification, scheduling and optimization of charging behavior. The first step involves the implementation of clustering algorithms (k-means and DBSCAN) on consumer load profiles to determine EV users according to unique energy consumption behavior. Then, time-series forecasting models are used, such as LSTM, ARIMA, and Prophet, to forecast the variation of short-term loads and find the best off-peak charging hours. A First-Come, First-Served (FCFS) scheduling system is then used to assign charging slots to make sure that users are treated fairly, transparently, and equally to gain access to the grid. A rescheduling dynamic module reacts to sudden load spikes and automatically reallocates charging sessions, informing users via a smart web/mobile application interface. The simulation outcomes are supposed to prove that the suggested system will be able to identify the EV users with an accuracy of more than 90 per cent, decrease the peak load stress greatly, and increase the transformer lifespan. In addition, the interactive application will encourage the participation of the user, which will encourage responsible use of electricity and increase trust in the operations of utility services. This architecture is also entirely compatible with the SCADA infrastructure and the new smart grid infrastructure of LESCO and promotes scalability and modularity to utilise it

nationwide. In general, the study can be valuable to the energy transition in Pakistan as it provides a predictive, user-friendly, and policy-consistent DSM solution that does not compromise grid stability but comforts the consumers. The framework will provide an international model of national EV integration, which will result in sustainable electrification, smart grid operation, and data-driven energy governance in developing economies.

Chapter 1: Introduction

1.1 Background

The transportation sector is amid a revolution as countries move from vehicles powered by internal combustion engines to electric vehicles (EVs) because of imperatives for climate change and energy security. By 2030, the global fleet of EVs is estimated at 350 million, with EVs accounting for around 60% of all new car sales globally (IEA, 2023). This transport electrification is one of the biggest changes in energy use since the industrial revolution, which could drastically change load profiles that electricity distribution networks have to cater for.

Pakistan has also seen the importance of this shift, and in 2019, the National Electric Vehicle Policy (NEVP) was introduced, setting ambitious goals for EV adoption. The policy also targets that by 2030, 30% of all new vehicle sales will be electric; it is estimated that there will be around 500,000 EVs on the Pakistani roads by 2030 as per the policy (Ministry of Climate Change, 2019). These targets put Pakistan in the same league with the rest of the world on decarbonization, but they also introduce new and unfamiliar challenges for the energy distribution network, which was designed and built decades before the widespread concept of residential EV charging.

Lahore Electric Supply Company (LESCO), which supplies electricity to more than 6.5 million users of Lahore and adjoining districts, is especially hard hit by the problem of managing the expected surge in demand. LESCO's network consists of around 146 grid stations, equipped with 10/13 MVA (132/11 kV) power transformers, and around 47,000 downstream distribution transformers, many of which are already running at full capacity in the peak summer months when air-conditioning loads are high (LESCO, 2023). The grid-station power transformer is the strategic asset of interest in this study, as it serves the demand of thousands of consumers, and its heightened ageing or failure under continuous overload has network-wide consequences, including replacement costs and lead times, which are far greater than those of a single distribution transformer. It is important to note that transformers are very sensitive to the addition of new EV loads to the existing load, as described by Gong et al. (2012), and that this burden passes through the transformer's windings to the grid-station power transformer, causing hot-spot temperatures and insulation ageing to exceed design ratings (Hilshey et al., 2013).

A key complicating factor for EV integration is that there is a fundamental information gap: utilities generally do not know which homes have EVs. EV purchases are not subject to notification by utilities, and EVs add to the total load profile of the house. This hiddenness hinders proactive grid planning. The disaggregation of these loads can be solved by the concept of Non-Intrusive Load.

Monitoring (NILM), which was first proposed by Hart (1992). The early NILM research was based on a simple on/off appliance, but in recent years, with the development of machine learning, particularly unsupervised clustering algorithms like k-means and DBSCAN, there are promising methods to identify complex shape EV signatures in smart meter data (Wang et al., 2019; Parson et al., 2012).

Demand-Side Management (DSM) is becoming a more cost-effective approach to accommodating new loads on existing grid systems, thereby avoiding costly additions to the system. But, for effective DSM, it is necessary to have a reliable identification of controllable loads and prediction of the grid conditions. Deep learning techniques, in particular Long Short-Term Memory (LSTM) networks, have shown better results in electrical load forecasting (Kong et al., 2019), which can be used to build up intelligent systems that can suggest optimal charging windows according to the predicted grid state.

1.2 Problem Statement

With the expected increase in EV operation as per NEVP 2019, LESCO's network faces a significant risk to its integrity of operation and asset life. The primary technical issue is transformer overloading at the grid station, where the total loading of the feeders to the station, combined with simultaneous charging from EVs, can cause the transformer to be overloaded, resulting in an excessive hotspot temperature in the windings, accelerated winding insulation degradation and eventual transformer failure. Overloading is a much more serious failure mode of a single 10/13 MVA power transformer, since it supplies thousands of consumers, compared to any single distribution transformer.

LESCO's grid-station power transformers have a natural cooling (ONAN) rating of 10MVA and can be forced-air cooled (ONAF) to 13MVA, but on hot summer evenings, they are already being used at 80–90% of their ONAN rating, with AC load dominating and little margin available for new EV load. At night, when hundreds of Level 2 EVs are charging, each EV uses 7-22 kW, which could overload the power transformer if they all happen to be on the same feeders at the same time. Research has shown that even a low level of EV penetration (15 to 20%) can cause 30 to 40 per cent higher peak loading for transformers, which is beyond their thermal capability and can speed up transformer insulation ageing (Shafiee et al., 2013). The 50-500kVA distribution transformers, which are each subject to the same superposition effect at the feeder level, are analyzed as a secondary, distribution-level validation case in this thesis.

Transformer overloading is more than just capacity issues. The relationship between ageing of transformer insulation and operating temperature is exponential: the higher the temperature, the shorter the insulation life (IEEE Std C57.91, 2011) - every 6-8°C above the rated temperature, the insulation life is reduced by 50%. Thermal stress accumulates and overloading of the transformer for short periods of time again degrades the insulation and causes dielectric failure. The cost of replacement of a failed distribution transformer for LESCO is on the order of PKR 1.5-3 million, without the cost-of-service interruption and compensation to the customers. There are 47,000 transformers in service, and if there are even a few failures, it is a significant financial risk.

One of the main challenges in EV charging load management is that the utility can't determine who has EVs in their home. EV adoption is not done in cooperation with the utility, but rather vehicles are sold at dealerships and charging equipment is installed by private electricians, without any information being sent to the distribution company. This implicitness is a drawback for LESCO: it is unable to demand to manage programmed, unable to see the level of EV penetration at the transformer, and unable to proactively reinforce feeders in areas with high EV penetration. The utility has no way to know where EVs are connected, and so it can only work in a reactive manner and only after the events of transformer overload or several customer complaints start to mount up.

There are two inadequate solutions to EV grid integration. Hardware solutions, such as transformer upgrades, dedicated EV charging feeders, and smart charging infrastructure, will come at a cost of PKR 50-100 million per grid station, which is unattainable for a developing country's utility company with competing infrastructure requirements. Simple Time of Use (ToU) tariffs apply to simple software solutions and do not consider the dynamic nature of grid conditions, and are unable to adapt to the spatial heterogeneity of EV adoption on LESCO's service territory.

An intelligent, software-based demand-side management system is needed to: (1) classify EV-owning households from consumption data using machine learning techniques, (2) predict the short-term load condition of the transformer, and (3) optimize the charging schedule to avoid transformer overload and meet the user demands. It should be a system that would work within the existing infrastructure of LESCO and would not require any hardware upgrades; it must give actionable intelligence for grid management.

1.3 Research Aim

The main objective of this research is to design an Artificial Intelligence-based Demand-Side Management (AI-DSM) framework for LESCO to reduce the extra load on the 10/13 MVA grid-station power transformer due to charging of EVs, the framework classifies EV-owning

households in consumption patterns using clustering algorithms, predict short-term loading on the power transformer using a deep learning algorithm, and optimizes the charging schedule to ensure the loading on the transformer remains within the IEEE C57.91 thermal limitations of the transformer to prevent accelerated ageing of the transformer insulation while meeting user convenience requirements.

1.4 Research Objectives

This study is guided by the following objectives:

1. To create and test an EV user identification system to identify households with EVs in smart meter data using unsupervised machine learning algorithms (DBSCAN and K-means clustering) to provide for targeted demand management interventions.
2. To build and test a short-term load forecasting model for forecasting the loading of the 10/13 MVA grid station power transformer for 24 hours ahead using a Long Short-Term Memory (LSTM) network, which will comprise time-series features and weather parameters that will be useful in Lahore's cooling load-dominated scenario.
3. To develop and apply a charging-schedule optimization algorithm to identify an optimal charging window for EVs, minimizing the loading on the power transformer during the peak period, subject to the transformer's capacity and winding hot-spot temperature constraint, user departure deadline and minimum charging requirement.

1.5 Research Questions

The aim of this thesis is to answer the following research questions:

1. To what extent can EV-owning households be accurately distinguished from household load profiles using unsupervised clustering algorithms (DBSCAN and k-means), and what features in household consumption offer the highest discriminatory power in detecting EVs?
2. What level of forecasting accuracy can the LSTM-based models achieve for the loading on LESCO's 10/13 MVA grid-station power transformer, and how do the results compare with those of the baseline statistical models in different seasons?
3. What is the potential reduction in peak loading of the 10/13 MVA power transformer and capacity/thermal-limit violation for different EV penetration levels for the integrated AI-DSM framework, compared with the uncoordinated charging framework?

1.6 Scope of the Study

The study is on residential charging coordination in LESCO's network, where the main grid constraint is the 10/13 MVA Grid- Station Power Transformer, and the secondary case is the 200 kVA distribution transformer, which is located at the feeder level. The level 2 AC charging (7-22KW) is used as the benchmark for residential charging, as this is the most common charging mode internationally and is in line with the developing EV charging infrastructure in Pakistan.

AI-DSM is a combined technical framework of three components. The EV identification module uses DBSCAN and k-means clustering on load profiles from households to distinguish EV-owning households from non-EV-owning households. The load forecasting module is based on LSTM neural networks to forecast the energy demand of transformers for 24-hour horizons. The charging optimization module is a constrained optimization based on transformer capacity for calculating the optimal charging windows.

The study area is geographically defined as LESCO's service area, and the methodology is designed in such a way that it can be replicated by other DCOs in Pakistan. The simulation environment has load patterns typical of the residential/grid-station loads in Pakistan and includes seasonal variations and weekly cycles.

The research utilizes simulation-based validation, for which load data sets are generated based on LESCO consumption data, and publicly available smart-meter data is adapted to the Pakistani context. This approach has the disadvantage that it does not allow for any claims on real-life implementation, but it allows for the controlled application of the experiments for EV penetration scenarios that would be difficult to implement in Pakistan at the moment.

1.7 Significance of Study

This research makes several contributions to the academic and practical aspects of utility planning.

From a theoretical point of view, the study contributes to the use of unsupervised learning in non-intrusive load monitoring in developing country scenarios where the smart meter penetration is not high, and ground labels are not available. A clustering-based EV identification approach, combined with deep learning forecasting and mathematical optimization, offers a holistic end-to-end solution for the entire EV management process, from identification to scheduling.

In terms of practicality, the AI-DSM framework could triple LESCO's ability to manage data related to its current operations. First, EV identification can help with proactive grid planning by showing spatial patterns of EV adoption prior to transformer failures. Secondly, the reliability of

Load forecasting aids in decision-making for operation, load dispatch and contingency planning. Third, with optimized charging recommendations, a mechanism is created to postpone expensive upgrades to transformers without compromising service reliability.

The transformer-centric approach directly looks at the most urgent asset management issue faced by LESCO. The framework prevents overloading of transformers and helps prolong their service life, lowers failure rates and postpones capital investment in new transformers. For a utility with 47,000 distribution transformers, even small savings in useful life can add up to significant savings.

The broader energy transition for Pakistan can benefit from this research, as it can know that grid-side hurdles in EV adoption can be resolved by intelligent software solutions instead of very costly infrastructure investment.

1.8 Definition of Key Terms

Artificial Intelligence-Driven Demand-Side Management (AI-DSM): A load management technique that uses machine learning algorithms to identify EV users, predict grid conditions, and optimize charging times to ensure system efficiency, without overloading distribution assets.

An unsupervised machine learning technique where data points are grouped by similarity measures without labelled training data. DBSCAN (Density-Based Spatial Clustering of Applications with Noise) and k-means are two popular methods that make different assumptions regarding the geometry of clusters.

Power-Transformer Overloading: When the total load of a grid-station power transformer (10/13 MVA, 132/11 kV) is greater than its rated load, it increases the hot-spot temperature of the windings beyond the design limits, which means an increase in insulation ageing. This is the main failure mode considered in this study; the same fault mode at the distribution-transformer (50–500 kVA) level is considered as a secondary case.

Electric Vehicle Penetration Level: Per cent of households in the service area of a distribution transformer that own and regularly charge an EV, which represents the total EV charging load to be coordinated.

Load Forecasting: The prediction of the future demand of electricity from past electrical consumption and explanatory variables. Day-ahead charging schedule optimization is possible by short-term forecasting based on a 24-hour horizon.

Long Short-Term Memory (LSTM) Network: A type of RNN that can learn long-term temporal dependencies in sequential data, suitable for electrical load prediction for detecting complicated consumption patterns.

Optimal Charging Window: The time slot where an EV can charge in order to minimize its contribution to transformer peak load and maintain a charge level that reaches the driver's desired departure time.

Transformer Capacity Constraint: The maximum amount of permissible loading, which in this study was assumed to be 90% of the power transformer's ONAF rating (10% safety margin), that the charging optimization needs to adhere to, to prevent the winding hot-spot temperature from exceeding the IEEE C57.91 limit, and to ensure that the insulation does not fail prematurely.

1.9 Thesis Outline

This thesis consists of six chapters as follows:

Chapter 1: Introduction sets the context of the research in the Pakistani EV policy arena and EV transformer overloading problems of LESCO; sets research aim, objectives and questions; and sets the scope and significance of the study.

Chapter 2: Literature Review provides a critical analysis of the literature regarding EV load identification through clustering methods, short-term load forecasting methods with a focus on deep learning approaches and charging coordination optimization strategies. The chapter presents the research gaps of this study.

Chapter 3: Methodology outlines the research methodology, comprising data sources, data preprocessing, implementation of DBSCAN and k-means for EV identification, LSTM architecture and training for load forecasting, problem formulation of the optimization, and evaluation metrics.

Chapter 4: Results and Analysis contains experimental results, such as the clustering performance for EV detection, accuracy comparisons of the forecasting models, optimization results under the penetration scenarios and the integrated framework results.

Chapter 5: Discussion presents findings in relation to the stated objectives, explores implications for LESCO's operations, comments on the findings in relation to other literature and acknowledges limitations.

Chapter 2: Literature Review

The shift towards an electrified transportation is a paradigm shift for power distribution systems and requires a revisit of traditional grid planning and operation strategies. This chapter reviews the current state of the literature on EV impact on distribution assets, Non-Intrusive Load Identification (NIL) in poor data environments, and the development of short-term load forecasting (STLF) methods. The review reveals that there is a great lack of literature on the use of unsupervised machine learning for residential EV management in the context of developing grids with no “ground truth” data and very limited infrastructure options.

2.1 The Impact of EVs on Distribution Networks

The addition of EVs represents the first introduction of high-power, stochastic loads in distribution networks designed for predictable, passive domestic loads. From the early adoption stages, researchers such as Clement-Nyns et al. (2010) have proven that uncoordinated charging of Plug-in Hybrid Electric Vehicles (PHEVs) causes substantial power losses and voltage deviations which are not manageable by conventional protection schemes. It is the distribution transformer that is cited repeatedly as the major obstacle in this transition, at least for the existing residential infrastructure, where investments in upgrading the infrastructure are slow.

The life of a distribution transformer depends mainly on the thermal degradation of its winding insulation made of cellulose paper and mineral oil. The operating temperature-insulation life relationship is determined by the Arrhenius reaction rate theory as outlined in the IEEE Guide for Loading Mineral-Oil-Immersed Transformers (IEEE Std C57.91, 2011). This standard is based on the theory that there is a logarithmic relationship between the rate of ageing acceleration factor ($FAAF_{AA}$) and the winding hottest-spot temperature. There is a critical threshold above which the rate of ageing will change rapidly, normally above 110°C for thermally upgraded papers. A well-known statement based on this model, which also agrees with recent empirical analysis by Grey and Morsi (2015), is that for every 6°C that the operating temperature exceeds the temperature rating, the insulation life is effectively cut in half. This non-linearity implies that brief periods of excessive overload can have a disproportionately high effect on this property, while long periods of light loading will not have the same effect.

Both Hilshey et al. (2013) and Gong et al. (2012) quantified this impact related to EV charging. Hilshey's probabilistic model showed how uncoordinated charging can cause equivalent ageing.

The factor of transformers is expected to increase by orders of magnitude, and in high penetration scenarios, the equivalent age of a 40-year transformer can be less than 10 years. Gong et al. (2012) pointed out that a few EVs, if charged at the same time in hot-spot situations, can cause the temperature to exceed the critical level. This thermal stress is aggravated by power quality problems; Deilami et al. (2011) demonstrated that random charging can cause a significant degradation of the voltage profile and an increase in losses in the power system, which causes more thermal stress to the transformer core.

A very important behaviour factor that is usually ignored in technical studies only is the "clustering effect", Muratori (2018). Based on a simulation of real-world residential feeders, Muratori found that EV adoption isn't evenly distributed across the country, but rather is spread out in certain affluent neighbourhoods, because of social influence and peer effects. That means that although system-wide EV penetration is low (e.g., 5%), local EV penetration may be 50-70% (Richardson et al., 2012). The resulting concentration forms "hotspots" of failure that are not taken into account at the system level during planning, resulting in unplanned outages.

Geographic bias in literature - Contextual Divide: A critical synthesis. Most of the studies (e.g., Clement-Nyns et al., 2010; Deilami et al., 2011) make the assumptions of "Smart Grid" with advanced metering infrastructure (AMI), strong communication backbones, and well-defined network topologies. On the other hand, there is little literature about the development of contexts such as Pakistan. Developing grids with "blind" operations, without real-time monitoring below the feeder level, and where assets are already highly loaded, because of the rapid urbanization (LESCO, 2023). Many of the proposed solutions do not apply to utilities such as LESCO without extensive adaptation because they assume the availability of the data required for grid management, which is not the case in the West.

2.2 EV Load Identification Techniques (NILM)

The crucial first step for the utility to successfully manage EV loads is to first be able to detect them in the aggregate household load. This challenge is related to Non-Intrusive Load Monitoring (NILM) as introduced by Hart (1992). Identification of the individual appliances by their steady state change of real power and reactive power was proposed by Hart. Modern EV charging has more complex signatures that may overlap with other high-power resistive loads, such as HVAC systems, whereas toasters or incandescent bulbs are simple resistive loads.

Early work, including that of Parson et al. (2012), takes a probabilistic approach to disaggregate loads using probabilistic models and Hidden Markov Models (HMMs). These methods needed a

A lot of knowledge of appliance cycles and could not cope with the varying power consumption of the modern inverter-based appliances. Recently, it has been turned into Deep Learning. To this end, Zhang et al. (2020) developed a Sequence-to-Point (Seq2Point) learning framework, where the load disaggregation task was modelled as a sequence translation task like a language translation task, and achieved state-of-the-art accuracy. Similar high-precision socio-demographic identification was developed in Wang et al. (2019) to classify certain household characteristics from smart meter data.

But these supervised models are fundamentally limited by the lack of very large, very frequent, labelled datasets in which the occurrence of an EV's "ground truth" (i.e., when the EV was charging) is known; customers don't spontaneously report EV purchases, and collection of such "ground truth" through sub-meter installation is too expensive for a utility like LESCO. This means that we need to move towards unsupervised learning.

A more popular algorithm is the k-means clustering, which is fast but can only find spherical clusters of similar variances – a condition that is often not met by the variation in load shapes. Unsupervised clustering provides a way to detect EVs based on distinct shapes or changes of intensity in the daily load profiles when using lower frequency data. Theoretically superior are density-based methods such as DBSCAN (Density-Based Spatial Clustering of Applications with Noise). Unlike other methods, DBSCAN can detect clusters of arbitrary shapes and, importantly, detect "noise" (outliers). Because EV charging profiles are often very high-power use, concentrated in a relatively small amount of time, and distinct from the variable, "baseload" power use of a household, clustering based on density shows promise. Cominola et al. (2017) discussed unsupervised methods and noted that they typically don't match supervised methods in absolute accuracy, but they provide the scalability needed to deploy in the real world, where labels are not abundant.

2.3 Short-Term Load Forecasting (STLF)

Having identified the EV loads, the grid operator should be able to forecast the future load states in order to schedule charging in an efficient way. Traditional forecasting in the utility industry was based on linear time series techniques such as ARIMA (Auto-Regressive Integrated Moving Average) and SARIMA. These models are good for aggregated, smooth demand at a regional level, but they do not work as well with the high fluctuations and non-linear relationships caused by residential charging of EVs, say Raza and Khosravi (2015). Its stochastic nature causes a "bursty" and non-stationary load profile—thus, the fundamental assumptions of linear regression are

broken. The LSTM Revolution Today, the field has turned solidly in the direction of Artificial Intelligence, or Recurrent Neural Networks (RNNs). The seminal work of Hochreiter and Schmidhuber (1997) on Long Short-Term Memory (LSTM) networks gave the architecture that was necessary to address the "vanishing gradient" problem and help networks learn long-term temporal dependencies. In the realm of the smart grid, Zheng et al. (2017) showed that LSTMs can be used to model such long-term dependencies (e.g. how the temperature yesterday affects the load today) and can outperform traditional forecasting methods. This was also supported by Kong et al. (2019), who validated the models for households, revealing that LSTMs are much better at capturing "bursty" household demand as compared to Feed-Forward Neural Networks (FFNNs).

Hybrid Approaches, Local Context. Recent progress has been directed towards hybrid and deep architectures, such as Marino et al. (2016), who stacked LSTMs for building energy forecasting, which has led to improved accuracy. A persistent gap is, however, that most international studies are trained on US/EU data and limit the study only to the distribution/building level, with a very limited integration of machine-learning forecasting and coordination to thermal protection of the upstream grid-station power transformer in resource-constrained networks. A key focus of this research is closing that gap, that is, bridging the data-driven management of EVs and the loss of life of the 10/13 MVA power transformer.

2.4 Summary and Research Gap

It is verified by literature that EVs are a serious thermal risk for transformers (Muratori, 2018; Hilshey et al., 2013), and LSTM networks are the best prediction model (Kong et al., 2019). But there is still a disconnection in the Identification phase in the development areas. Rigorous comparison of DBSCAN and k-means for strictly residential EV detection in unsupervised cases of no ground truth labels. Moreover, not much research sets out to combine all three- unsupervised detection, LSTM forecasting and optimization into a single hardware-free framework, especially for a resource-constrained utility such as LESCO. The purpose of this research is to fill this void.

Chapter 3: Methodology

The research design, data sources, and algorithmic implementations supporting the proposed AI-Driven Demand-Side Management (AI-DSM) framework are described in this chapter. The methodology was designed in such a way that it could be reproduced for future researchers and utility practitioners to replicate and adapt for their situations. The chapter is divided into five components: research design and conceptual framework, data collection and preprocessing, detailed explanation of three main technical modules (EV User Identification, Load Forecasting, and Charging Optimization), and conclusions. All sections are detailed enough to be replicated, and include hyperparameter selections, software libraries, and evaluation metrics.

3.1 Research Design and Framework Overview

The research is quantitative and experimental with a simulation design. Since the EV market is still in a very early stage in Pakistan and no real-world EV charging data exists for the country publicly available from LESCO, reconstructed load data obtained from the consumption statistics published by LESCO is used, in addition to smart-meter data in the international context, adapted to the Pakistan scenario. This may constrain claims of readiness for use in the real world, but will allow for controlled testing of varying percentages of EVs in the fleet that would not be possible at today's EV adoption levels. This is in line with the approaches for foundational EV-grid integration studies (Hilshey et al., 2013; Muratori, 2018).

3.1.2 Research Philosophy and Approach

The research is based upon positivism, where observable phenomena such as load patterns and transformer temperature are assumed to be measurable and modifiable. The research methodology is mainly deductive, i.e., from theory to fact: theories on thermal ageing of transformers (IEEE Std C57.91) and theory on machine learning (LSTM Networks) are applied to the context of LESCO's distribution network. The experimental approach is to vary the levels of the EVs (the penetration level of EVs in the system) and the clustering algorithm type and observe their corresponding impact on the dependent variables (classification accuracy, forecasting error, peak load reduction).

Modularity of the proposed AI-DSM framework: The proposed AI-DSM framework is modular, consisting of three components in a sequence, which represent an end-to-end process from household-level data consumption to actionable charging recommendations. Modularity also guarantees that every component can be tested, improved or exchanged when more efficient




Figure 3. 1: AI-DSM Framework Conceptual Diagram

algorithms are discovered. Such a design is important for deployment: A utility could replace the forecasting module, for instance; in this case, they might move from LSTM to a newer Transformer-based model, without needing to change the identification or optimization modules.

The framework functions as follows: (1) The extracted features from the smart meter data are clustered based on clustering algorithms, and the households are classified into EV-owning and non-EV-owning as illustrated in Figure 3.1. In the Forecasting module, historical load data, weather variables, and calendar features are fed into the module to forecast the total load on a distribution transformer for the next 24 hours. (3) The Optimization module is responsible for setting the optimum time slots for charging the EVs based on the predicted load profile that will cause less stress to the transformer while satisfying the energy requirements of users. The data generated by each module is directly passed on to the next, ensuring a smooth flow of information.

This applies to the Grid Station or Distribution Transformer.

There are important scoping decisions to be made as to the hierarchy of the network to which this framework is applicable. LESCO's network consists of grid station transformers of 10/13MVA (132/11kV) with 10MVA normal rating and 13MVA emergency rating under enhanced ONAN/ONAF cooling and secondary distribution transformers (SDT) of 50-500kVA (11kV/400V) serving residential neighborhoods. In keeping with the original research proposal, the analysis reported in this thesis focuses on the grid-station power transformer (the 10/13 MVA, 132/11 kV transformer), and a secondary, feeder-level reference case is the 200 kVA distribution transformer that supplies 80 residential households.

There are three reasons for this scoping decision. First, the grid-station power transformer represents the load from a whole feeder group of residential customers; thus, coordinating EV charging at this level avoids being loaded by EVs on the single most capital-intensive asset in the local grid network, and provides an added insight into the system-wide diversity of loads that is not visible at the feeder level. Secondly, when EV adoption increases, the synchronous evening load of thousands of EVs peaks at the grid-station transformer, which is the last thermal choke

point, after the feeder level headroom has been exhausted. Third, since each grid-station problem can be seen as an aggregation of distribution-transformer sub-problems, the algorithms developed at the grid-station level (identification, forecasting, and optimization) are directly transferable down the hierarchy to LESCO's distribution transformers, of which the utility operates approximately 47,000 (compared to 196 grid stations).

Level	Asset	Rating	Voltage	Role in This Study
Grid Station	Power transformer	10/13MVA (normal/emergency)	132/11 kV	Primary focus: binding grid constraint for EV loading
Distribution	Distribution transformer	200 kVA	11 kV/400 V	Secondary, feeder-level reference case
Household	Service connection	~5–10 kW typical peak	400/230 V	Data source for clustering and aggregation

Table 3. 1: Network Hierarchy and Scope

3.2 Data Collection and Preprocessing

Three different data streams are combined in the data pipeline for this research: Residential load profiles, Meteorological data, and EV charging templates. This section provides information on the source, nature and processing of the streams.

Residential and Grid-Station Load Profiles. Data privacy regulations and commercial sensitivity result in non-public availability of granular smart-meter data from LESCO. This study, therefore, develops load profiles for the residential and grid-station sectors, consistent with the overall consumption pattern and statistics reported in LESCO's Annual Report, 2022-2023. Building of the profiles is done in a bottom-up manner, by aggregating a typical residential load curve, which is based on the typical duty cycles of various appliances, like lighting, fans, refrigerator, air conditioner, cooking, etc., weighted by the number of appliances present in the households of Lahore, to the feeder level and then to the grid-station level. The resulting profiles show the typical overnight peak (18:00-22:00) observed in LESCO, consumption is in line with the reported per consumer annual value, and the granularity of data is at 15 minutes, which is like the smart meters being deployed through the AMI rollout in Pakistan. This method guarantees the repeatability of the research and minimizes potential limitations in access to data while still capturing the reality of LESCO's documented load features.

Weather data (Lahore Hourly Weather Data): Historical weather data, such as Hourly temperature, Hourly humidity, and Solar Irradiance, is obtained from the free and open-source weather API, Open-Meteo. Since subsistence air conditioning load in Lahore's subtropical climate is strongly associated with the temperature, it is considered a major exogenous variable for the forecasting model, as the temperatures usually cross 40°C in the summer. Humidity is taken up as a secondary variable because it has a bearing on the feeling of comfort and, consequently, the use of air conditioning. Solar irradiance is added to model daytime load shapes, even though the contribution of solar to residential load is relatively small without significant penetration of solar PV on the roof.

EV charging profiles and EV charging Templates are modelled on the basis of the Level 2 AC charging specification. A typical session is a constant 7kW draw (32A, 220V, single phase) for the length of time based on battery size and State of Charge (SoC).

Parameter	Value	Justification
Charger Power (Level 2)	7 kW	Standard home charger; 32A @ 220V
Battery Capacity	60 kWh	Typical for mid-range EVs (e.g., Hyundai Kona, Tesla Model 3)
Starting SoC	30%	Represents a moderate daily commute (~50 km)
Target SoC	90%	Preserves battery longevity; common charging target
Charging Duration	~5.1 hours	$(0.9 - 0.3) \times 60 \text{ kWh} / 7 \text{ kW} = 5.14 \text{ hrs}$
Arrival Distribution	$\mu = 19:00$, $\sigma = 1.5$ hours	Post-work arrival, consistent with Muratori (2018)
Departure Distribution	$\mu = 07:00$, $\sigma = 1.0$ hours	Morning commute departure
Charging Probability	85% per day	Not every day requires charging

Table 3. 2: EV Charging Session Parameters

EV Load Injection. For the 30% EV penetration level used in this study, Level 2 charging signals are overlaid on 24 randomly selected household profiles from the 80 baseline household profiles, based on the probabilistic arrival model in Table 3.2. As the set of EV-charging households is

known by construction, this renders a labelled dataset for assessing the unsupervised clustering algorithms.

Parameter	Value
Total Households	80
EV Households	24 (30% penetration)
Baseline Households	56 (70%)
Study Period	October 2022 – September 2023 (12 months)
Data Resolution	15 minutes
Total Transformer-Level Records	34,945
Total Household-Level Records	2,795,600
Weather Variables	Temperature (°C), Humidity (%), Solar Irradiance (W/m ²)
Calendar Features	Day-of-week, Weekend flag, Holiday flag, Ramadan flag
Ramadan Period	23 March – 21 April 2023
Public Holidays Encoded	15 dates (Pakistan calendar 2022–2023)

Table 3. 3: Dataset Summary

3.2.1 Data Preprocessing Pipeline

The raw data is preprocessed to guarantee quality and compatibility with the machine learning models in several steps:

The first stage involves linear interpolation, where the gap is less than an hour, and forward fill, where the gap is greater than one hour, and then a rolling mean smoothing over one hour. Observations that are extreme outliers (out of range) are restricted to within 3σ of the time-of-day mean. The values of all the loads are scaled between 0 and 1 for training the neural network (Min-Max scaling). Categorical day-of-week features are one-hot encoded into 7 binary features, whilst the binary public-holiday flag is kept as is.

3.2.2 Data Partitioning Strategy

The entire dataset is split into three chronologically ordered parts for training the model,

selecting hyperparameters, and evaluating the performance of the model without bias:

Partition	Proportion	Date Range	Records	Purpose
Training	70%	1 Oct 2022 – 12 Jun 2023	24,461	Model parameter estimation
Validation	15%	12 Jun 2023 – 6 Aug 2023	5,242	Early stopping; hyperparameter tuning
Testing	15%	6 Aug 2023 – 30 Sep 2023	5,242	Out-of-sample performance evaluation

Table 3. 4: Data Partitioning

The chronological split is essential to avoid temporal data leakage, that is, not to include future observations in the training set. The validation set is used in two ways: To monitor early stopping in LSTM training (patience = 10 epochs on validation loss) and to choose the best hyperparameters. The results presented in Chapter 4 are all on the held-out test set, which is not used in training any of the models to give an unbiased estimate of the generalization performance.

3.3 Module 1: EV User Identification

The goal of this module is to divide the households into two categories: EV owners and non-EV owners. In a real-world deployment, data is not labelled, so two unsupervised clustering algorithms, k-means and DBSCAN, are implemented and compared.

The daily load profiles for each household are aggregated, and the top ten statistical features are extracted for the entire period of observation spanning one year (12 months). The features are chosen because they are theoretically relevant to the physical aspects of Level 2 charging and can capture the unique temporal and magnitude signatures of a 7-kW charging event on top of typical Pakistani residential consumption.

Feature	Description	Rationale
evening_mean (kW)	Mean power, 18:00–24:00	EV charging occurs predominantly during evening hours

evening_std (kW)	Standard deviation, 18:00–24:00	Captures variability introduced by intermittent charging
night_mean (kW)	Mean power, 00:00–06:00	Control feature; not expected to discriminate
daily_max_mean (kW)	Mean of daily peak loads	EV charging elevates daily maxima substantially
daily_range (kW)	Mean of (daily max – min)	A wider range indicates the presence of large intermittent loads
peak_to_trough_ratio	Mean of daily max / daily min	A higher ratio indicates “peaky” profiles
ramp_rate_max (kW)	Mean of max	ΔP
load_factor	Mean load/peak load	Lower values indicate peakier profiles from EV charging
evening_energy_share	Σ evening / Σ total energy	A higher share suggests evening-concentrated charging
total_daily_energy (kWh)	Mean daily energy consumption	EV households consume substantially more energy

Table 3. 5: Features Extracted for Clustering

To avoid dominance of distance calculation by features with greater magnitude, all features are z-score normalized to zero mean and unit variance before clustering.

3.3.1 k-means Implementation

The k-means algorithm is set with $k = 2$ (EV and non-EV clusters), k-means++ initialisation, Euclidean distance, convergence tolerance value of 10^{-4} and a maximum of 300 iterations. Implementation is done with Python and Scikit-learn.

3.3.2 DBSCAN Implementation

DBSCAN will have two hyperparameters: 1) ϵ : radius of the neighbourhood, and 2)

MinPts: Minimum number of points that will make up a core. Based on the convention of Ester et al. (1996), minPts is set to $2 \times \text{dimensionality} = 2 \times 10 = 20$. Sorting the k-th nearest-neighbour distances (with $k = \text{MinPts}$), plotting and finding the elbow point defines the ϵ value. Scikit-learn is used for implementation with Euclidean distance.

3.3.3 Validation Metrics

Since ground truth labels are available, the standard classification metrics are used for performance evaluation, namely:

Metric	Formula	Interpretation
Precision	$TP / (TP + FP)$	Of all flagged EV households, what fraction actually are?
Recall	$TP / (TP + FN)$	Of all actual EV households, what fraction were identified?
F1-Score	$2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$	Harmonic mean balances both
Silhouette Score	$(b - a) / \max(a, b)$	Cluster separation quality; range $[-1, +1]$
Cohen's d	$(\mu_1 - \mu_2) / s_{\text{pooled}}$	Effect size quantifying group separation per feature

Table 3. 6: Clustering Evaluation Metrics

3.4 Module 2: Load Forecasting

The aggregate apparent-power loading of the 10/13 MVA grid-station power transformer is here predicted for the following 24 hours (96-time steps at 15-minute resolution). The Optimization module is informed of periods of high and low demand thanks to accurate prediction.

3.4.1 Model Architecture

A stacked Long Short-Term Memory (LSTM) network is used, which has been proven to be superior in capturing long-term temporal dependencies in sequential data (Kong et al., 2019; Zheng et al., 2017). The architecture is composed of two LSTM layers (with hidden

units 64 and 32) and a Dense output layer. After every LSTM layer, dropout regularization (rate = 0.2) is performed.

Hyperparameter	Value	Justification
LSTM Layers	2 (stacked)	Captures multi-scale temporal patterns
Hidden Units	64, 32	Sufficient capacity for residential loads; prevents overfitting
Dropout Rate	0.2	Standard regularization reduces variance
Lookback Window	96 steps (24 hours)	Captures full daily cycle
Forecast Horizon	1 step ahead	Single-step prediction for maximum accuracy
Optimizer	Adam (lr = 0.001)	Adaptive learning rate; robust to hyperparameter choice
Loss Function	MSE	Standard regression loss; penalizes large errors
Epochs	100 (Early Stopping)	Patience = 10 on validation loss
Batch Size	32	Balances gradient stability with memory

Table 3. 7: LSTM Model Hyperparameters

3.4.2 Input Feature Vector

The input features at each time step consist of the lagged value of the load as well as exogenous variables such as temperature (°C), humidity (%), day-of-week (one-hot, 7 features) and a public-holiday flag (binary). The exogenous drivers are dominated by temperature, indicating the cooling-dominated nature of the residential and the grid-station demands in Lahore and its high correlation with the summer evening peak.

3.4.3 Training Protocol

The data is divided into a training set (70%), validation set (15%) and test set (15%) as indicated in Table 3.4 chronologically. The model is trained by the Adam optimizer.

Early stopping uses tracking of the validation loss with patience 10. Implementation is based on TensorFlow/Keras in Python.

3.4.4 Baseline Model

For comparison, a Seasonal Naive baseline is implemented, predicting the value for each time step as the value observed one week earlier at the same time of the year. This is the bare-minimum viable forecasting approach without using machine learning.

Forecasting Evaluation Metrics

Metric	Formula	Interpretation
MAE	$(1/n) \sum$	$y_i - \hat{y}_i$
RMSE	$\sqrt{[(1/n) \sum (y_i - \hat{y}_i)^2]}$	Penalizes large errors; same units as MAE
MAPE	$(100/n) \sum$	$y_i - \hat{y}_i$
R²	$1 - SS_{res} / SS_{tot}$	Variance explained; 1 = perfect

Table 3. 8: Forecasting Evaluation Metrics

3.5 Module 3: Charging Optimization

This module aims to find the optimal charging schedule for all the identified EV households connected to a single distribution transformer, while keeping the total load below the transformer's rated load capacity in peak hours. Three charging strategies are tested so that a comprehensive evaluation of the coordination value is possible.

3.5.1 Charging Scenarios

Scenario	Description	Coordination Level	Requirements
Uncoordinated	EVs charge immediately upon arrival at full power	None	No infrastructure needed
FCFS	EVs queue by arrival time; charging pauses when the load reaches transformer	Reactive	Real-time load monitoring

	capacity		
AI-DSM Coordinated	Charging scheduled into forecast valley periods via quadratic optimization	Predictive	Load forecast + optimizer

Table 3. 9: Charging Scenario Definitions

The use of the First-Come-First-Served (FCFS) scenario as an intermediate benchmark allows a detailed evaluation of the added value of predictive optimization compared to reactive coordination.

3.5.2 Problem Formulation (AI-DSM)

This optimization problem is expressed in the form of a Quadratic Program (QP). Define $P_{base}(t)$ as the predicted baseline load at time t (as in Module 2) and let $P_{charger, i}$ be the rated charging power of EV i . The decision variable is $x_i(t) \in [0, 1]$, denoting the amount of charging power that EV i uses at time t when it is fully charged.

Objective Function: minimise the sum of square aggregate loads over the scheduling horizon:

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

$T = 96$ (24 hours with 15 min resolution) and N is the number of identified EV households. This squared formulation has the effect of creating a penalty for peak loading, relative to valley loading, and thus “flattening” the load curve.

3.5.3 Constraints:

Optimization is subject to (1) Capacity constraint with a 10% safety margin – the total load should not exceed 90% of the transformer's rated capacity (e.g., 10/13 MVA power transformer on its ONAF rating) during any time step to allow a margin of error for forecasting and to keep the winding hot-spot temperature within IEEE C57.91 limits. (2) Energy requirement — ALL of the required energy must be provided to each EV before

departure. (3) Charging window — charging only takes place between the arrival and departure times. (4) Variable bounds — $0 \leq x_i(t) \leq 1$.

Component	Specification
Problem Type	Quadratic Program (QP)
Decision Variables	$x_i(t) \in [0, 1]$ for all EVs i and time steps t
Objective	Minimized sum of squared total load
Constraints	Capacity (with 10% safety margin), Energy, Window, Bounds
Safety Margin	10% — effective ceiling = $0.9 \times P_{\text{rated}}$
Solver	OSQP (via CVXPY)
Input	Predicted baseline $P_{\text{base}}(t)$, EV arrival/departure, energy requirements
Output	Optimal charging schedule $x_i^*(t)$

Table 3. 10: Optimization Problem Summary

3.5.4 Transformer Ageing Assessment

The IEEE Standard C57.91-2011 ageing-acceleration factor is used to evaluate the ageing of transformer insulation:

$$F_{\text{AA}} = \exp(15000/383 - 15000 / (\theta_{\text{HS}} + 273))$$

Instead of a static approximation, the IEEE C57.91-2011 dynamic top-oil/hot-spot thermal model described in Section 3.6 is used to determine the winding hot-spot temperature (θ_{HS}) in °C. The value of 1.0 is the rated ageing rate ($\theta_{\text{HS}} = 110$ °C); values greater than 1.0 are accelerated ageing. The daily loss-of-life is found by summing up the area under the 24-hour profile of F_{AA} . This allows for a straight comparison of the long-term impact on infrastructure of the three charging scenarios on the 10/13 MVA power transformer (Sections 4.3 and 4.4).

3.6 Power-Transformer Thermal Modelling (IEEE C57.91-2011)

Since the research target is to reduce the impact on the 10/13 MVA power transformer of the grid-station during EV loading, this section details the dynamic thermal model, which is used to convert the transformer loading into the hot-spot temperature of the windings, age acceleration, and loss-of-life. The model is based on the IEEE Guide for Loading Mineral-Oil-Immersed

Transformers (IEEE Std C57.91-2011, Clause 7) and models the dominant thermal path from ambient to top oil to winding hotspot, consisting of two first-order responses with different time constants.

The grid-station load is reconstructed with 15-minute resolution for 12 months and aligned with the feeder statistics data of LESCO Annual Report 2022-2023 and Lahore's meteorological data. The diversified evening peak of ~9.1 MVA is about 91% of the 10 MVA ONAN rating, and it closely matches what LESCO reported for their grid transformers in summer (LESCO, 2023), leaving the transformer dangerously vulnerable to further EV load.

With K the per-unit loading referenced to the 13 MVA ONAF rating, R the ratio of rated load loss to no-load loss, and n the oil exponent, the ultimate top-oil temperature rise over ambient is:

$$\Delta\theta_{TO,ult} = \Delta\theta_{TO,R} \cdot [(K^2 \cdot R + 1) / (R + 1)] ^ n$$

When the top-oil temperature gets closer to this value, with the oil time constant τ_{TO} , over each time step, Δt :

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

The rise of the hot-spot over top-oil is $\Delta\theta_{H, ult} = \Delta\theta_{H, R} \cdot K^{(2m)}$, and it responds with the much shorter winding time constant τ_w . The winding hot-spot temperature is the sum of the ambient, top-oil-rise and hot-spot-gradient components:

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

Table 3.11 lists the thermal parameters adopted for the 10/13 MVA, 132/11 kV ONAN/ONAF power transformer, consistent with typical 65 °C-rise mineral-oil units in IEEE Std C57.91-2011.

Parameter	Value
ONAN / ONAF rating	10 / 13 MVA
Rated top-oil rise over ambient ($\Delta\theta_{TO, R}$)	55 °C
Rated hot-spot rise over top-oil ($\Delta\theta_{H, R}$)	25 °C
Load-loss / no-load-loss ratio (R)	8
Oil exponent (n)	0.9
Winding exponent (m)	0.8
Top-oil time constant (τ_{TO})	150 min
Winding time constant (τ_w)	7 min
Power factor	0.95
Normal insulation life	180,000 h

Table 3. 11: IEEE C57.91 thermal parameters for the 10/13 MVA power transformer

The hot-spot temperature is used to calculate the ageing-acceleration factor $F_{AA} = \exp(15000/383$

– $15000/(\theta_{HS} + 273)$) and the percentage loss-of-life in a horizon is given by the integration of F_{AA} with respect to the nominal insulation life of 180,000 hours. This dynamic formulation is used in the loading, optimization and thermal analyses of Sections 4.3 and 4.4 for the 10/13 MVA power transformer.

Chapter 4: Results and Analysis

The empirical results of the AI-DSM implementation on LESCO's network are presented in this chapter. To highlight the data-driven basis of the framework, EV User Identification using unsupervised clustering (Section 4.1) and Short-Term Load Forecasting using LSTM networks (Section 4.2) are performed; to emphasize the coordination logic at a fine (feeder) level, Charging Optimization at distribution transformer level (Section 4.3) is presented; and to quantify the aggregation of the EV burden, the resulting winding hot-spot temperature rise and the loss-of-life implications under the IEEE C57.91 model, the central contribution of this study, thermal-impact analysis on the 10/13 MVA grid-station power transformer (Section 4.4), is presented. The integrated pipeline is assessed in section 4.5. Each module provides quantitative scores, graphical analysis and interpretive comments based on Chapter 3.

Forecasting and charging optimisation were carried out on the 10/13 MVA (132/11 kV) grid-station power transformer, with the 12-month load record reconstructed at 15-minute resolution and calibrated with the consumption statistics published in the LESCO Annual Report 2022-2023. The EV-identification module (Section 4.1), which is a household-level evaluation, unrelated to transformer rating, was validated and applied to a representative sample of 80 residential households within the transformer's catchment area, with 24 households being EV-owning based on the signature of Level 2 home-charging. The load record was divided into 70% of the data for training, 15% for validation and 15% for out-of-sample testing, to avoid temporal data leakage and reflect the real-life situation where forecasting models have to generalise to future, unseen periods.

4.1 EV User Identification Performance

The first research objective was to determine if the unsupervised clustering algorithms could be used to successfully classify EV-owning households from the larger population of residential customers based on aggregate load profile features only, without the need for labelled training data or self-reporting by customers. One of the two algorithms, k-means and DBSCAN, was implemented on a standardised feature matrix, and the results were compared to the known household classification.

4.1.1 Dataset and Feature Extraction

The clustering analysis was performed on statistical characteristics of the daily loads of each household computed over the entire one-year period. Ten features were calculated and chosen for being theoretically relevant to the physical nature of Level 2 residential charging, considering the

literature search (Section 2.2). The feature set includes temporal concentration (evening mean, evening standard deviation, evening energy share), magnitude signatures (daily maximum, daily range, maximum to minimum ratio), dynamics (maximum ramp rate), efficiency shape (load factor) and total consumption (daily energy, overnight mean).

Parameter	Value
Total Households	80
EV Households	24 (30%)
Baseline Households	56 (70%)
Analysis Period	October 2022 – September 2023 (12 months)
Total Load Records	2,795,600 (15-minute resolution)
Features Extracted	10

Table 4. 1: Dataset Summary for Clustering Analysis

Table 4.2 presents descriptive statistics for each feature stratified by household type, accompanied by inferential test results quantifying group separation.

Feature	Baseline Mean (SD)	EV Mean (SD)	Cohen's <i>d</i>	<i>p</i> - value	Discrimination
evening_mean (kW)	0.740 (0.267)	2.254 (0.157)	6.49	<0.001	Very High
evening_std (kW)	0.261 (0.045)	0.576 (0.072)	5.49	<0.001	Very High
night_mean (kW)	0.439 (0.197)	0.477 (0.097)	0.22	0.437	Low
daily_max_mean (kW)	1.396 (0.463)	4.100 (0.286)	6.87	<0.001	Very High
daily_range (kW)	1.278 (0.432)	3.996 (0.266)	7.69	<0.001	Very High
peak_to_trough_ratio	13.38 (7.55)	50.83 (18.26)	3.06	<0.001	Very High
ramp_rate_max (kW)	0.859 (0.138)	3.171 (0.225)	12.96	<0.001	Very High
load_factor	0.428 (0.020)	0.278 (0.032)	−5.88	<0.001	Very High

evening_energy_share	0.385 (0.015)	0.649 (0.057)	6.71	<0.001	Very High
total_daily_energy (kWh)	11.59 (4.57)	21.17 (3.36)	2.54	<0.001	Very High

Table 4. 2: Feature Statistics by Household Type

Note: All p-values are from independent samples Welch's t-tests. Cohen's $d > 0.8$ is conventionally considered a large effect size.

Some patterns should be closely examined. The feature with the most discriminative power is `ramp_rate_max`, a Cohen's d of 12.96, which represents an extraordinarily large effect size and corresponds to the steep load ramp-up when a 7 kW charger is activated, whereas the residential energy profile averages to below 2 kW. In the Pakistani context, no other common household appliance offers a step change of a similar magnitude and duration. The next largest appliance load in the home is the air conditioning unit (AC), and these have a compressor modulation cycle and rarely exceed 2.5 kW per unit (Ahmad et al., 2018), giving a ramp rate which falls within the baseline distribution. The `daily_range` and `daily_max_mean` features also capture the amplitude signature of EV charging, with effect sizes greater than 6.8, which means that they are essentially non-overlapping distributions between the two groups.

One of the features did not discriminate between household types: `night_mean` (Cohen's $d = 0.22$, $p = 0.437$). This negative finding is instructive, not discouraging. The nights are dominated by base load, which is not systematically influenced by EV ownership, such as refrigeration, standby electronics and intermittent fan operation. The discriminative signal is in the evening period, when post-commute charging is in session, and not in the overnight period, when charging may or may not be in session, depending on the length of the charging session and when the car arrives. The time of day that this signal occurs has important consequences for feature engineering in deployment: features from the evening window are crucial, while overnight statistics can probably be omitted without any loss of classification performance, which decreases the amount of computation needed for large-scale screening.

Because it is in the opposite direction to the magnitude features, the `load_factor` feature merits a separate discussion. The mean load factor of EV households (0.278) is lower than that of the baseline households (0.428). A low load factor is the "peakier" signature – the average consumption is smaller as a fraction of the peak – and this is exactly the effect of intermittent high power charging events overlaid on a lower baseload. That inverse relationship was posited in the

methodology. (Section 3.3) because of Muratori's (2018) description of EV charging as creating sharp demand peaks in otherwise relatively flat residential demand profiles, which is very evident from the data.

4.1.2 Clustering Algorithm Performance

Both algorithms were used on the standardized (zero mean, unit variance) feature matrix with k-means parameters $k = 2$ clusters, k-means++ initialization, Euclidean distance and convergence tolerance of 10^{-4} . Based on the k-distance graph method, the hyperparameters of DBSCAN were set as follows: $\text{MinPts} = 2 \times 10 = 20$ (where 10 is the feature dimensionality), and the elbow of the sorted k-th nearest-neighbour distance graph gave the optimal value of $\epsilon = 4.368$.

Metric	k-means	DBSCAN
Precision	1.000	1.000
Recall	1.000	1.000
F1-Score	1.000	1.000
Silhouette Score	0.744	0.744
Noise Points	0	0

Table 4. 3: Clustering Algorithm Performance Comparison

Both algorithms were able to be perfectly classified, with all 24 EV households classified as EV, and all of the 56 baseline households classified as non-EV, without false positives or false negatives. Both have the same confusion matrix:

	Predicted: Baseline	Predicted: EV
Actual: Baseline	56	0
Actual: EV	0	24

Table 4. 4: Confusion Matrix (Both Algorithms)

Maximizing the success of classification by two very different algorithmic techniques: a centroid-based method and a density-based method, requires explanation, not just celebration. The feature statistics in Table 4.2 are rethought to make sense of the result. Cohen's d for the rate of ramp is 12.96, which means that the distributions for the two groups are almost 13 pooled standard deviations apart. Even with the worst-case scenario for the distribution, the chance that any given data point will be misclassified as left or right by a midpoint decision boundary is miniscule. This can be quantified using the silhouette score of 0.744 – values greater than 0.7 indicate a good

cluster structure – the two groups are in separate, well-separated regions of the ten-dimensional feature space.

Perfect recall (1.000) is no trivial matter. In the EV identification context, a false negative occurs when a charging household goes undetected and doesn't get coordinated, thus allowing the charging load to continue to cause peak stress. Both algorithms can include all EV households at the 30% level of penetration, and these feature magnitudes. It is an important question whether this performance can be maintained with lower penetration or with smaller charging devices (e.g. 3.3 kW) to be explored in the future, because the signal-to-noise ratio would be significantly reduced.

4.1.3 Dimensionality Reduction Visualization

To confirm the quantitative results, the standardized feature space was reduced to two dimensions using PCA and t-SNE.




Figure 4. 1: presents the PCA projection in three panels: (a) ground truth labels, (b) k-means predictions, and (c) DBSCAN predictions.

Eigen analysis of the first two principal components accounts for 98.6% of the total variance (PC1: 80.7%, PC2: 18.0%). Because of the dominance of PC1, it is important to interpret the result. Standardization of input features means that 80.7% of the variance is concentrated on a single axis, meaning that the main cause of inter-household variance is the correlated set of features related to EV charging – evening mean, ramp rate, daily max – which are all highly represented in this component. EV households form a compact upper-right quadrant, which is clearly distinct from the more diffuse baseline quadrant, and for which there are no data points that connect each quadrant. It is clear from the visual impression that the two populations are not statistically different but are separate in the reduced space. A practical implication is that even a single univariate threshold on PC1 alone would achieve almost 100% classification.




Figure 4. 2: presents the t-SNE projection, which preserves local neighborhood structure and provides a complementary nonlinear view

The t-SNE embedding validates the PCA results and provides a slight twist which the linear projection cannot provide: the points within each cluster are cohesive. EV households do not just appear to be further from baseline households, but rather each EV household is closer to the other EV households than to any baseline household. This internal density is contributing to the theoretical appropriateness of density-based techniques (e.g. DBSCAN) that are dependent on local density and not on global centroids. The same performance between these two algorithmic families is due to the excellent separability of the data and not a constraint imposed by either algorithm.

4.1.1 Feature Importance Analysis




Figure 4. 3: presents box plots comparing the distribution of the six most discriminative features between the two household types.

The visual analysis complements the quantitative results in Table 4.2, which cannot be done with the results from the summary statistics. With `ramp_rate_max`, 100% of the EV interquartile range is above the 99th percentile of the baseline distribution, which means there is literally no overlap between the two populations. Similar but slightly less separation is found in the `evening_mean` and `daily_max_mean` features. By contrast, the `peak-to-trough ratio` also has a high discriminative power ($d = 3.06$), but is much more spread in the EV group, with some EV households having ratios around 25, and others having ratios that exceed 80. This variation within groups is due to the charging frequency, as a household that charges every day at the same time will have extreme ratios, and a household that charges sporadically will have intermediate ratios. The practical consequence is that if only one of the features is considered, then it is possible that the irregular chargers will be misclassified; the multivariate approach that is taken here allows these different profiles to be captured by the complementary information provided by the various features.

4.1.4 Sensitivity Analysis

Robustness was evaluated by systematically corrupting the feature matrix with additive Gaussian noise from 0 to $\sigma=1.0$ to model various sources of noise and meter inaccuracies, imputation artefacts and atypical household behaviour.




Figure 4. 4: plots F1-Score against noise level for both algorithms.

Noise σ	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
k-means	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.96
F1	0	0	0	0	0	0	0	0	0	0	0
DBSCA	1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N F1	0	0	0	0	0	0	0	0	0	0	0

Table 4. 5: Classification Performance Under Feature Noise

Both algorithms have drastically different robustness profiles, which should be carefully analyzed: k-means is able to maintain perfect or near-perfect classification over the whole range of noise levels, with a minimum F1 of 0.960 at $\sigma = 1.0$. This is because of k-means' global optimization of centroid positions: symmetric noise in the feature distributions affects individual data points, but the centroids stay in place due to the overall feature distributions. The inter-cluster gap found in the

PCA analysis is sufficiently large so that the noise needed to "connect" clusters is higher than the point where the raw meter data becomes essentially unusable.

In contrast, DBSCAN completely failed for 0.1. This is the mechanism of failure; when the noise is making the data points more dispersed, the data points in the fixed ϵ -neighborhood (chosen based on clean data) are not enough to meet the $\text{MinPts} = 20$ condition, and thus the algorithm determines that most or all the points are noise. Note that this is not an inherent drawback of density-based methods, but rather a result of taking a fixed ϵ with varying noise. Adaptive ϵ -selection using an estimated data quality, e.g., scaling ϵ with the estimated noise level, might recover performance. However, this extra complexity is a drawback to the simplicity of DBSCAN in practice, which makes it suitable for exploration analysis.

The divergent robustness profiles make a strong statement: In practice, k-means is the much more reliable option when deploying it at LESCO, where data quality of smart meters might differ between substations and seasons. DBSCAN is sensitive to the ϵ parameter and may be useful in an analytical environment where data quality can be checked, but not automated in an automated, unsupervised screening environment on a large scale.

4.1.5 Summary of EV Identification Results

The experimental assessments lead to three main conclusions. First, at the 30% EV penetration level modelled in this study, both k-means and DBSCAN perfectly classify the 24 EV households, without any false positives. Second, the discriminative signal is dominated by features capturing the magnitude and temporal concentration of charging – maximum ramp rate ($d = 12.96$), daily range ($d = 7.69$), and daily maximum ($d = 6.87$) – and these features had extreme effect sizes, which is why the separability was exceptional. Third, k-means is recommended for deployment due to its good noise robustness ($F1 \geq 0.96$ for all noise levels tested), simple algorithm (only hyperparameter is known a priori, $k = 2$ is from binary classification objective) and efficient computation ($O(nk)$ vs $O(n \log n)$ for DBSCAN).

The findings for the first research question in Chapter 1, "Can unsupervised clustering detect EV-owning households from residential load profiles with high accuracy?" are that it is indeed possible, and that the most discriminative features include the evening mean consumption, maximum ramp rate and daily peak demand.

4.2 Load Forecasting Accuracy

The second research objective involved the development and validation of a short-term load forecasting model that could accurately forecast the loading at the transformer level, thereby enabling the charging optimization for the next day. This section compares the LSTM-based model with the Seasonal Naive baseline model in terms of their prediction accuracy, error characteristics, and temperature sensitivity.

4.2.1 Experimental Configuration

Forecasting experiments focused on the apparent-power loading for the 10/13 MVA grid-station power transformer and were based on the 12-month load record reconstructed at 15-minute resolution (Section 3.2). The model is a stacked Long Short-Term Memory (LSTM) network with two recurrent layers of 64 and 32 units and 0.1 dropout, which takes a 96-step (24-hour) history of the load as an input and extracts the next 96 quarter-hourly load values by utilizing the day-ahead weather and calendar covariates (temperature, humidity, time of day and day of week) for the forecast horizon. It was trained with mean-squared-error loss using the Adam optimizer and early stopping on a held-out validation set.

Partition	Records	Proportion	Period
Training	24,595	70%	Jul 2023 - Mar 2024
Validation	5,270	15%	Mar 2024 - May 2024
Test	5,271	15%	May 2024 - Jun 2024

Table 4. 6: Data Partition for Forecasting

The Seasonal Naive baseline uses the actual observation from the same time on the previous day as a prediction for each time step, a simple, transparent baseline that would be the minimum viable forecasting strategy a utility could use without the machine-learning capability. Daily weather variations and weekday and weekend loading variations in residential areas are significant in Lahore, and the forecast for the same time yesterday will present a significant systematic error to compare against the learned model.

4.2.2 Model Performance Comparison

Table 4.7 presents the quantitative comparison, evaluated on the held-out test set.

Metric	Seasonal Naive	LSTM	Improvement
MAE (MVA)	0.363	0.168	53.8%
RMSE (MVA)	0.526	0.229	56.4%
MAPE (%)	10.13	4.86	52.0%
R-squared	0.861	0.974	-

Table 4. 7: Forecasting Performance Metrics

The changes are significant. The LSTM's MAE is 0.168 MVA, which is around 1.3% of the transformer's ONAF (13 MVA), and 5.9% of the mean loading of the transformer, a level identified by Kong et al. (2019) as sufficient for operational day-ahead scheduling. The R-squared value is 0.974, which indicates that the model accounts for 97.4% of the variance in load during the test period, with the remainder representing the stochastic fluctuations in demand, which are not predictable from the aggregate weather and calendar signals.

It is noteworthy that the magnitude of the improvement of MAE (53.8%) indicates that the Seasonal Naive baseline, which captures the average of the daily shape, cannot account for the day-to-day variability engendered by the unpredictable summer weather in Lahore and by the weekday/weekend differences. The LSTM learns the nonlinear relationship between ambient temperature and air-conditioning load and leverages a day-ahead weather forecast on the horizon to adapt to the fluctuations and to achieve a much lower error than the 35-40% MAPE reductions proposed by Kong et al. (2019) for air-conditioning load forecasting using LSTM models on the residential side.

The improvement in RMSE (56.4%) is greater than the improvement in MAE, which suggests that LSTM is good at limiting large prediction errors. This is operationally critical as the RMSE will penalise larger forecast errors quadratically, which can lead to either over-curtailling EV charging when the forecast is too high, or unsafe transformer loading when the forecast is too low. The LSTM's tight error distribution significantly decreases the likelihood of both failure modes.

4.2.3 Forecast Tracking Analysis




Figure 4. 5: Illustrates the LSTM's ability to track actual load over a representative week in the test period, compared with the Seasonal Naive prediction.

Several characteristic behaviours can be observed in visual analysis. The LSTM model also reproduces the double peak that is found in the residential demand, namely the morning peak around the activity period and the evening peak during the cooking/entertainment activity, with good accuracy during the week. Evening peak magnitudes are also reproduced with very little systematic error – a property which is important for the downstream optimization module, which needs to evaluate the available capacity headroom for EV scheduling. The shaded $\pm 8\%$ confidence band includes almost all the observations, indicating that the implicit uncertainty in the model is well contained for planning purposes.

The Seasonal Naive prediction follows the overall pattern of the load curve but shows significant phase and amplitude errors, especially during the transition between weekdays and weekends when there are changes in the activity pattern. Such tracking failures would be translated into the optimization module, resulting in either unnecessarily reducing EV charging (by overestimating the baseline) or accepting EV charging that leads to a genuine overloading (by underestimating the baseline). This close monitoring by LSTM reduces the margin of error for the optimizer to work within, allowing it to make use of even more of the capacity available without endangering safety.

4.2.4 Error Distribution Analysis




Figure 4. 6: presents the error distribution comparison and a normal probability (Q-Q) plot for the LSTM residuals.

The LSTM's error distribution is concentrated around zero with a standard deviation of around 0.23 MVA, which is less than half of the Seasonal Naive spread of 0.53 MVA. The Q-Q plot indicates that the residuals are approximately normal, except that the large quantiles show some heavy tailing, which may include some larger errors when sudden changes in the weather occur. These tail events are associated with sudden changes in cooling demands, which are stochastic and cannot be adequately reduced through time-series modelling alone.

The approximation of the error distribution by a normal distribution is used to support the use of parametric safety margins in the optimization module. A plus-minus 2-sigma load (plus-minus 0.46 MVA) means that there is approximately 95% confidence that the true load is in the band, allowing the optimiser to use a principled (as opposed to ad hoc) capacity derating.

4.2.5 Temperature Sensitivity

The subtropical climate of Lahore, characterized by high summer temperatures of more than 45°C, leads to a high temperature–load coupling which is driven by the air-conditioning demand. It is therefore important to study the behaviour of LSTM under thermal stress.




Figure 4. 7: stratifies the absolute forecast error by temperature range across the test period.

The model results have similar accuracy throughout the complete range of temperatures, and the median absolute errors have not been too high even in the extremely high temperatures (above 40 C). At higher temperatures, there is no material degradation: as the temperature is an exogenous input, the LSTM uses the temperature-load relationship as another predictive signal, thus effectively turning a source of volatility into a source of predictability.

4.2.6 Summary of Forecasting Results

On held-out test data, the LSTM has an R-squared of 0.974, outperforming the Seasonal Naive baseline by 53.8% (MAE), 56.4% (RMSE) and 52.0% (MAPE). It is generalised over the day-to-day weather fluctuations of the period tested and is consistent throughout the entire temperature range of Lahore. The forecasts supply the basic load predictions that the charging-optimisation module is based on, and which consequently limit the effectiveness of the downstream scheduling.

4.3 Charging Optimization on the 10/13 MVA Power Transformer

The third objective of the research was to develop and test a charging optimization strategy to avoid transformer overload and meet EV user needs. Three different scheduling strategies are compared: uncoordinated charging (a natural baseline), First-Come-First-Served (FCFS) queue-based

scheduling, and AI-DSM optimized scheduling. Adding FCFS as an intermediate strategy will give supervisory observation a sense of the less complex coordination benchmark needed to explain the importance of full optimization.

4.3.1 Simulation Configuration

Optimization experiments are simulated on a typical summer peak day on the 10/13 MVA (132/11 kV) grid-station power transformer with approximately 4,000 residential connections. This equates to 1,200 electric vehicles, each taking 7 kW from a Level 2 AC charger, for a mean daily energy demand of 20 kWh per vehicle (23.9 MWh in total). On this single worst day, the three coordination strategies are compared to the no-EV scenario.

Parameter	Value	Basis
Grid-station transformer	10/13 MVA ONAN/ONAF, 132/11 kV	Study subject
Residential connections	about 4,000	Grid-station catchment
EV penetration	30% (1,200 EVs)	Modelled adoption level
Charger rating	7 kW (Level 2 AC)	Typical home charging
Mean energy per EV	20 kWh (SD 5)	Daily commute replenishment
Total daily EV energy	23.9 MWh	1,200 EVs in aggregate
Mean plug-in/departure	18:55 / 07:00	Hardman et al. (2018)
AI-DSM safety cap	90% of 13 MVA ONAF (11.7 MVA)	10% forecast-error margin
Representative day	Summer peak, ambient 47 C	Worst-case loading

Table 4. 8: Optimization Simulation Parameters




Table 4. 9: presents the distribution of EV plug-in arrival times across 1,200 vehicles.

Most arrivals occur between 17:00 and 21:00 with a mean of 18:55, which is typical of the post-work home arrivals reported by Hardman et al. (2018). The average departure time of 06:49 allows for a buffer of ~12 hours overnight, which is well above the ~4 hours charging time. This time, Slack is the resource that coordination strategies are exploiting—otherwise, there would be no other time slots available, and no scheduling approach could do better than uncoordinated charging.

4.3.2 Scenario Definitions

The charging infrastructure does not provide sufficient information for electric vehicles. EV owners connect to the power upon arrival and charge at maximum capacity until the energy need is satisfied. No communication takes place with the utility. This is the natural baseline (i.e., not as a result of any DSM intervention) and the worst-case scenario for transformer stress, as the charging onset is timed to coincide with the pre-existing residential evening peak.

The 'first-come, first-served' (FCFS) scenario. Individual charging is suspended if the overall charging load of EVs will exceed the rating of the transformer (13MVA ONAF), which gets loaded in sequence by arrival time. The next queue EV runs when capacity is available: when baseline load decreases, or previous EV sessions are finished. FCFS only needs real-time monitoring of the load and a queue mechanism, no load forecasting, no optimization computation, and no advance scheduling.

Because of its simplicity, it is a practical measure of the extent to which coordination is necessary and is a reasonable benchmark for the least effective form of coordination possible.

Scenario C - AI-DSM Coordinated. The framework runs the quadratic programming formulation described in Section 3.5 to determine all the EV charging time in the valley periods identified by the 24-hour-ahead forecast made by the LSTM (mainly at nighttime between 00:00 and 06:00). To account for forecast errors and stochastic load variation, there is a 10% provisioned buffer, reducing the actual ceiling to 90% of the ONAF rating of 13 MVA (11.7 MVA). Given the transformer capacity, energy requirement, and charging-window constraints, the optimizer distributes charging so as to minimize the sum of squared aggregate load in the 24-hour horizon.

4.3.3 Daily Load Profile Comparison




Table 4. 10: presents the transformer load profile under all three charging scenarios alongside the no-EV baseline and the 13 MVA ONAF capacity limit.

Load shapes are different, in a qualitative sense. When the charging is not coordinated with the others, the simultaneous occurrence of 1200 EV chargers and the natural evening peak results in a huge power spike between 18:00 and 23:00. During the evening peak, the total load is 15.57MVA, which is 119.8% of the 13MVA ONAF rating and 155.7% of the 10MVA ONAN rating, due to most households returning from work at the same time and connecting their appliances and also the base load of cooking and entertainment is at its peak. The FCFS strategy keeps the peak at the capacity limit (13.0 MVA, 100% ONAF) but does so at the expense of a more spread-out high-load time in the evening and into the night, as queued EVs take turns having access. The AI-DSM

A coordinated approach does exactly the opposite: it moves practically every EV load into the overnight valley between 00:00 and 06:00, and leaves the load profile in the evening, pretty much unchanged from the no-EV situation, and builds a new high plateau overnight, utilizing a previously unused part of the profile.

The difference between FCFS and AI-DSM is not just quantitative but qualitative. FCFS is reactive, meaning it reacts to the capacity violations at the time they are about to happen, while AI-DSM is predictive, which charges during times of known low demand. The reactive approach will have to work (or be close to working) at the constraint boundary in the evening, while the predictive approach will not have to work at all during the constrained period.

4.3.4 EV Charging Load Distribution




Table 4. 11: isolates the EV charging component to illustrate how each strategy distributes the 1,200-vehicle charging load across 24 hours.

Three different time trends can be distinguished. In this worst-case scenario, the uncoordinated EV load is displayed in Panel (a) as a concentrated block during the hours 18:00-22:00, which exactly coincides with the baseline peak. The EV component contributes approximately 5.6 MVA to the total of 15.57 MVA at the evening peak on top of the ~9.95 MVA baseline. As seen in panel (b), FCFS allows some evening charging within the capacity envelope, but due to the progressive discharge of queued vehicles, the majority are postponed to a long tail after midnight, and as a result, the evenings still include a significant number of EVs. The charging pattern for the AI-DSM is shown in panel (c), which clearly exhibits the clean valley-filling pattern, characterized by charging concentrated between 00:00 and 06:00, with levels subsequently distributed so as to have a near-constant aggregate load, which is a good indicator of successful load flattening.

4.3.5 Performance Metrics

Table 4.9 presents the comprehensive comparison across all four scenarios.

Metric	No EVs (Baseline)	Uncoordinated	FCFS	AI-DSM Coordinated
Peak load (MVA)	9.95	15.57	13.00	9.95
Peak loading (% ONAF)	76.5%	119.8%	100.0%	76.5%
Overload hours (>13 MVA)	0.0	2.5	0.0	0.0
Load factor (%)	51.7	39.9	47.6	62.2
Overnight valley avg (MVA)	3.44	3.44	3.44	5.52
Energy delivered (MWh)	-	24.3	23.7	23.9
User satisfaction (%)	-	100	99.3	100
Peak reduction vs uncoord.	-	-	16.5%	36.1%
Mean F_AA	0.009	0.256	0.111	0.012
Max F_AA	0.066	3.371	1.032	0.073
Daily loss-of-life (%)	0.00012	0.00341	0.00148	0.00016

Table 4. 12: Optimization Performance Comparison

Note: User Satisfaction is defined as the percentage of EVs receiving their full energy requirement by departure time.




Figure 4. 8: presents these metrics in bar chart form for visual comparison.

The findings need to be considered on a variety of levels.

Peak Reduction. The AI-DSM reduces the peak load from 15.57 MVA to 9.95 MVA, a 36.1% reduction, which causes the transformer to drop from an alarming 119.8% of the transformer's ONAF rating to a comfortable 76.5% of the rating. FCFS results in a 16.5% cut to precisely the 13 MVA ONAF rating (100.0%) with no margin for forecast errors, stochastic spikes, or unforeseen load increases. The difference between operating at 100% vs 76.5% of ONAF quantifies the value added by predictive optimisation, as compared to reactive queue management. In IEEE C57.91 terms, it is materially different hot-spot temperatures and ageing rates.

Overload Elimination. The transformer is subjected to uncoordinated charging that causes it to operate in overload (above 13 MVA ONAF rating) for 2.5 hours per day, which is sustained thermal stress, during which insulation deteriorates at a much higher rate than it's designed to withstand (according to the IEEE C57.91 exponential ageing model). FCFS and AI-DSM completely remove overloading. The mechanisms are different—FCFS is a reactive method that just pauses charging at the boundary, and AI-DSM is a proactive method that schedules charging away from the danger zone—but the result is the same—zero overhead hours. The difference is in the thermal profile in the following hours, as described in Section 4.3.6.

Load Factor. The AI-DSM achieves the highest load factor (62.2%), while the no-EV baseline and FCFS have a load factor of 51.7% and 47.6%, respectively, and uncoordinated charging achieves a load factor of 39.9%. This measures the efficiency of the use of the infrastructure: a higher load factor will mean more energy delivered per unit of transformer capacity installed. This improvement (39.9% to 62.2%) is purely a temporal redistribution of the same total energy, an increase of 56% in asset utilisation. The paradoxically low load factor of the uncoordinated scenario is due to the negative impact of peak superposition: When adding the EV energy solely during peak.

times, the denominator (peak load) increases at a higher rate than the numerator (average load), thus reducing the load factor.

Valley Utilisation. The most significant difference between FCFS and AI-DSM is the overnight valley loading, which is 5.52 MVA for AI-DSM and 3.44 MVA for FCFS. The overnight average is higher with the AI-DSM since all EV charging is pre-scheduled to the deepest valley by the AI-DSM, as opposed to FCFS, which only defers to overnight the vehicles it could not serve in the evening queue. This distinction is significant from a thermal point of view: IEEE C57.91 ambient corrections indicate that a particular load will generate less thermal stress overnight, when ambient temperatures are at the daily low, than in the hot evening, making the AI-DSM's ageing contribution even more significant than the load data would suggest.

Energy delivery and satisfaction. The coordinated strategies bring in nearly all the daily EV energy (AI-DSM 100% and FCFS 99.3%), resulting in almost all EVs being charged before the owner leaves. The coordination mechanism changes only the timing of energizing the energy without changing the amount of energy. This is important for user acceptance, because it does not require any mobility penalty, the overnight scheduling is clear to the owner and the car is plugged in on arrival, and the owner can use it on departure.

4.3.6 Transformer Loss-of-Life under the Dynamic Thermal Model

With the help of the dynamic thermal model of Section 3.6 of IEEE Standard C57.91-2011, the loading profile of each scenario and the ambient temperature during the day are transformed into a trajectory of the hot-spot of the windings and an ageing-acceleration factor:

$$F_{AA} = \exp\left(\frac{\theta_{HS} - 110}{6.5}\right)$$

An F_{AA} of 1.0 means that the ageing rate is the same as the rated ageing rate at the reference hotspot of 110 C, whereas ageing rates greater than or less than 1.0 mean accelerated or decelerated ageing, respectively. The dynamic model accounts for the thermal time constants of the oil and winding and the ambient temperature drop overnight, giving a much more representative measurement of the loss-of-life than a simple static approximation of the peak load. The mean and maximum values for F_{AA} , and the daily loss-of-life for each scenario are summarized in the bottom rows of Table 4.9.

The ageing analysis shows the most unambiguous quantitative evidence for coordinated charging and unveils the difference between FCFS and AI-DSM that is hidden by coarser metrics.

With uncoordinated charging, the hot-spot temperature is pushed far into the overload region in the evening, evolving the ageing-acceleration factor to a maximum of 3.37 and the daily loss-of-life to 0.00341% - roughly 28 times the no-EV baseline value of 0.00012%. When continued day after day on hot summer evenings, this is the predominant cause of premature insulation failure, as seen in the probabilistic analysis by Hilshey et al. (2013) and the empirical modelling by Gong et al. (2012), which both show that uncoordinated EV charging can have a significant impact on transformer life.

FCFS at 0.00148% loss-of-life per day will have a max F_{AA} of 1.03 in the extended near-capacity evening period, indicating that there will be some ageing slightly above the rated rate, but not of any significance daily, but rather over the multi-year time horizon relevant for asset management.

This is because the AI-DSM scenario is bound to a small fraction of the daily loss-of-life (0.00016%) and keeps F_{AA} at the maximum value of 0.073, as it focuses on charging during the cool valley hours overnight. In effect, it allows 30 % EV penetration with no measurable additional ageing of the grid-station transformer, which is of significant economic importance for LESCO as a 10/13 MVA transformer would cost them a few crores and have a procurement lead time of many years.

4.3.7 Comparative Assessment of FCFS versus AI-DSM

As a result, both FCFS and AI-DSM avoid transformer overloading and provide the total EV charging demands, which raises the question: what is the value of the complexity of AI-DSM (load forecasting, quadratic programming, and advanced scheduling) compared to the simplicity of FCFS (only real-time monitoring and a queue)?

The evidence suggests a differentiated reply in three aspects: First is the safety margin: while FCFS is operated at the edge of the envelope (100.0% ONAF peak), the AI-DSM operates within the envelope (76.5% ONAF peak), which means it's able to withstand stochastic load spikes without being taken down by them. For a real network with less than perfect metering and behaviour that is not exactly as modelled, operating without margin is precarious. Second, long-term asset health: the order of magnitude difference in mean F_{AA} (0.012 versus 0.111) and the difference in maximum F_{AA} (0.073 versus 1.032) suggest measurable differences in multi-year impact on transformer life. When magnified over a fleet of grid-station transformers over 20 years, the differences add up to significant capital-expenditure differentials. Third, infrastructure utilisation: The load factor of the AI-DSM (62.2%) is a significant improvement over FCFS

(47.6%) and indicates an improvement in the efficiency of asset-utilisation without any additional hardware.

Although neither of these methods will be seen as an operating issue in the short term, they both avoid the short-term acute problem of uncoordinated charging, and their collective long-term value over many years when applied to a large fleet of transformers is an important economic and engineering argument for the full optimization approach. FCFS can still be used as an interim solution in areas where computational infrastructure for predictive optimization is not yet in place, but it should be seen as a steppingstone and not an end goal.

4.3.8 Summary of Optimization Results

The AI-DSM framework is validated at the grid-station level through the optimization evaluation. Uncoordinated charging at 30% penetration causes the 10/13 MVA transformer to be at 119.8% of the transformer's ONAF for 2.5 hours per day, the hot spot ageing-acceleration factor to peak at 3.4 and the daily loss-of-life to increase to approximately 28 times the no-EV reference. FCFS scheduling does not overload but does run at the capacity limit with no safety margin, and the ageing is intermittent and is greater than the rated ageing. With the AI-DSM coordinated approach, the peak load is reduced by 36.1%, restores a comfortable margin (76.5% ONAF), achieves the highest load factor (62.2%) and keeps the ageing of the transformers near the baseline - while meeting essentially 100% of the user charging requirements without compromising the availability of vehicles.

4.4 Power Transformer (10/13 MVA) Thermal Impact Analysis

This section outlines the key findings of the study: a quantitative evaluation of the extra thermal load of EVs charging on the 10/13 MVA (132/11 kV) grid-station power transformer and how the AI-DSM framework reduces it. The loading profiles from section 4.3 are converted to winding temperatures and insulation loss-of-life (LoL) using the IEEE C57.91-2011 dynamic thermal model from section 3.6. A representative summer peak day (47 °C ambient max) is used, and the EV fleet is varied at 30% penetration of the approximately 4,000 residential connections served, along with a sweeping of the penetration from 0% to 50%.

4.4.1 From Loading to Winding Temperature

On a typical summer day, the diversified base load is already at 9.95 MVA - 99.5% of the transformer's 10 MVA ONAN rating and 76.5% of its 13 MVA ONAF rating - and the transformer.

It is already into its EV era with almost no natural headroom (Figure 4.9, Section 4.3). The peak is lifted to 15.57 MVA (155.7% ONAN / 119.8% ONAF) for 2.5 hours by the uncoordinated charging, is held at 13 MVA ONAF by FCFS and is returned to 9.95 MVA by the AI-DSM. The thermal impact of these very different loading patterns is then considered.

4.4.2 Winding Hot-Spot Temperature Response

The loading condition is translated to the winding temperature in the dynamic thermal model. Figure 4.12 shows the top-oil and hot-spot temperatures along with the ambient temperature as the evening load surge propagates through the oil thermal mass, into a delayed hot-spot peak that is more dramatic than the ambient temperature. A comparison of the hot-spot temperature for all the strategies is shown in Figure 4.13. The hot-spot peaks, at 85.2 C, on this 47 °C day with the no-EV baseline - hot for the day, but not quite up to the 110 C rated limit. Uncoordinated EV charging drives the peak hot-spot to 122.3 C - an increase of 37.1 C that takes the winding about 12 C beyond its thermal limit - while FCFS reaches 110.3 C, right at the limit. The AI-DSM strategy keeps the maximum hot-spot at 86.1 C, which is about the same as the without-EV baseline.




Figure 4. 9: Ambient, top-oil and winding hot-spot temperature evolution under uncoordinated EV charging (IEEE C57.91 dynamic model).




Figure 4. 10: Winding hot-spot temperature of the power transformer under each EV charging strategy. Uncoordinated charging exceeds the 110 C rated limit; AI-DSM holds it at the no-EV baseline.

Scenario	Peak (MVA)	% ONAN	% ONAF	Max Hot-Spot (C)	Hrs > 13 MVA	Daily LoL (%)
No-EV Baseline	9.95	99.5	76.5	85.2	0.0	0.00012
Uncoordinated	15.57	155.7	119.8	122.3	2.5	0.00341
FCFS	13.00	130.0	100.0	110.3	0.0	0.00148
AI-DSM Coordinated	9.95	99.5	76.5	86.1	0.0	0.00016

Table 4. 13: Loading and thermal outcomes on the 10/13 MVA power transformer (30% EV penetration) for all four scenarios.

4.4.3 Effect of EV Penetration on Temperature and Ageing

In order to characterise the burden as a function of adoption, EV penetration was swept from 0% to 50% for uncoordinated charging. The resulting peak hotspot temperature and daily loss-of-life are reported in Figure 4.14 and Table 4.11. The relationship is strongly non-linear: at 0%, the peak hot-spot temperature is 85.2 C; at 20%, it is 107.9 C; at 22% penetration, the peak hot-spot temperature is at the rated limit of 110 C; and at 50%, it is 153.4 C. The ageing of the insulation is an exponential function of the temperature, and the loss of life increases by more than two orders of

magnitude in this range, indicating that uncontrolled EV growth is a direct threat to the service life of the grid-station power transformer.




Figure 4. 11 Peak winding hot-spot temperature and daily loss-of-life of the power transformer versus EV penetration (uncoordinated charging).:

EV Penetration (%)	Peak Load (MVA)	Peak Loading (% ONAF)	Max Hot-Spot (C)	Max Top-Oil (C)	Max F_AA	Daily LoL (%)
0	9.95	76.5	85.2	69.3	0.07	0.00012
10	11.75	90.4	95.9	75.7	0.22	0.00030
20	13.62	104.8	107.9	83.5	0.81	0.00094
30	15.46	118.9	121.4	92.2	3.10	0.00329
40	17.37	133.6	136.9	102.2	13.05	0.01298
50	19.26	148.2	153.4	112.7	53.83	0.05017

Table 4. 14: Power-transformer loading, hot-spot temperature and loss-of-life versus EV penetration (uncoordinated charging).

4.4.4 Loss-of-Life and Mitigation

Figure 4.15 contrasts the daily insulation loss-of-life of the four strategies on a logarithmic scale. Relative to the no-EV baseline, uncoordinated charging raises the transformer's daily loss-of-life

by roughly 28-fold and drives the peak ageing-acceleration factor to 3.4 during the evening surge, whereas the AI-DSM strategy returns the daily loss-of-life to within a small fraction of the baseline. The thermal evidence, therefore, mirrors the loading evidence: intelligent coordination converts an EV burden that would otherwise consume disproportionate transformer life into one that the asset can absorb with no measurable additional ageing.




Figure 4. 12: Daily insulation loss-of-life of the 10/13 MVA power transformer by charging strategy (log scale, 30% EV penetration).

4.4.5 Summary of Power-Transformer Thermal Findings

At 30% EV penetration, the 10/13 MVA uncoordinated charging pushes the power transformer to 119.8% of the ONAF rating (155.7% of the ONAN rating) and increases the winding hot spot temperature by 37.1 °C to 122.3 °C, which is approximately 12 °C above the limit. When EV penetration reaches beyond approximately 22%, the limit is exceeded, and ageing accelerates exponentially. The AI-DSM framework ensures that the loading (76.5% ONAF) and hot-spot temperature (86.1 C) are kept at essentially pre-EV levels, while also providing 100% of the desired charging energy by scheduling charging into the valley time of night under a 90% capacity cap. These results directly answer the research aim, as they quantify the thermal burden on the grid-station power transformer due to EVs and show how this burden can be offset with software-based demand-side management (DSM), without any hardware reinforcement.

4.5 Integrated Framework Evaluation

The three modules (identification, forecasting and optimization) were planned as a "pipeline" in which each module follows the other. The overall performance should be evaluated as a whole, as incorrect results in the upstream modules will be carried over through the chain of interdependencies to the downstream modules.

The identification module was successful for the EV household roster: on the 80-household validation sample, it identified all 24 EV owners with no false positives or false negatives, which also scales to flag EV adopters among the grid station's ~4,000 connections. If a false negative (an EV house that is not caught) turned out to be the case, its charging would remain uncoordinated, thus risking a reintroduction of the stress peak the optimization is trying to avoid. At the feature separability demonstrated (Section 4.1), this risk is negligible at 30% penetration, increasing only at lower penetrations or at lower charger ratings, which should be quantified in future work.

The forecasting module provided the 24-hour ahead basic load forecast used by the optimizer to find valley periods and available capacity headroom. Forecast errors are small compared to the margins as indicated by an R-squared of 0.974 and an MAE of 0.168 MVA: a plus/minus 0.46 MVA (plus/minus 2-sigma) error band is within the 1.3 MVA difference between the 13 MVA ONAF rating and the 11.7 MVA ceiling of the optimizer. This forecast uncertainty was accounted for in the 10% safety margin incorporated in the optimization constraint.

The optimizer then synthesized both upstream inputs and came up with schedules that reduced the peaks by 36.1%, with no overload hours and no reduction in user satisfaction. The end-to-end performance from raw load data to actionable charging schedules shows that the three modules, although individually validated, work together as a complete system to provide results that would not be possible with any one module alone.

A deployed system would need some other components that are not modelled in this simulation: real-time monitoring for corrective feedback, communication infrastructure for disseminating the schedule, and exception-handling logic to be used in cases where there is a material mismatch between the actual and forecasted conditions (e.g., a significant outage on the feeder that was not planned for). These engineering requirements are beyond the scope of the present analytical study but are crucial steps along the path to a validated framework for operation.

Chapter 5: Discussion

This chapter brings all the experimental findings from Chapter 4 together and concludes their implications for the research questions from Chapter 1. These results are placed in the context of the existing literature on smart grid management, demand-side flexibility, and the difficulties of electricity distribution in developing economies. As Palensky and Dietrich (2011) explain, the paradigm shift from supply-following-demand to active coordination of flexible loads is especially relevant to the widespread deployment of distributed energy resources, such as electric vehicles, on distribution networks around the world.

The chapter is organized as follows: Section 5.1 provides an interpretation of the results from all three of the framework modules, incorporating results from the literature. The practical solution to the “blindness” problem, which motivated this research, is discussed in section 5.2. A phased implementation of LESCO is provided in Section 5.3. Section 5.4 assesses scalability. A critical appraisal of the limitations is presented in Section 5.5, and a comparison with previous research in the field is done in Section 5.6.

5.1 Interpretation of Findings

The experimental assessment of the AI-DSM framework showed that it was effective for all three research goals. Existing literature and the context of LESCO's distribution network are used to discuss each module below.

5.1.2 EV Identification Through Unsupervised Learning

The classification results of both k-means and DBSCAN at 30% EV penetration for the LESCO dataset were perfect ($F1 = 1.000$, Precision = 1.000, Recall = 1.000), with a DBSCAN silhouette score of 0.744, indicating a good cluster structure. These results significantly outperform identification accuracy reported in similar studies, such as Cominola et al. (2017) and Wang et al. (2018), which achieved F1 scores of 0.85–0.92 and classification accuracy of 87–91%, respectively, for disaggregation of appliances-level load in similar European residential settings, but with supervised learning methods requiring labelled training data which is not available for the LESCO case.

The outstanding performance can be explained by the feature-level analysis reported in Chapter 4. The most discriminative feature, maximum ramp rate, had a Cohen's d of 12.96, meaning that the separation of the EV and baseline distributions was almost 13 pooled standard deviations. This extreme separability is a direct result of the physical features of level 2 charging in the Pakistani

residential setting. The step change of 7 kW versus a baseline profile of 0.7-1.4 kW that a 7-kW charger provides cannot be mimicked by any other common household appliances. The next largest residential load in Lahore are the air conditioners, which cycle slowly, by compressor modulation, and which rarely use more than 2.5 kW per unit, meaning that the ramp rates are very well contained in the baseline distribution. The results agree with the seminal work of Hart (1992) on non-intrusive load monitoring, which showed that devices with high power consumption and a unique operating signature can be best disaggregated.

The PCA analysis showed that only 2 principal components are needed to account for 98.6% of the total variance of the inter-household scores, with the first component accounting for 80.7%. The discriminative information along this one axis suggests that the EV/charging binary split is the chief contributor to the variation in a household's characteristics that lies at the heart of why even basic clustering algorithms are able to perfectly separate the data. This separability would be less in more heterogeneous settings, where multiple high-power loads (heat pumps, pool heaters and home workshops) have similar signatures, and one would need more sophisticated approaches.

The practical insight provided by the sensitivity analysis was that k-means was more robust and less affected by feature noise, reaching $F1 \geq 0.96$ for all tested noise levels ($\sigma = 0-1.0$), whereas DBSCAN failed to be robust due to its sensitivity to the ϵ parameter and collapsed at $\sigma = 0.2$. This divergence comes with a very strong deployment recommendation. In operational screening at LESCO, the data quality of smart meters varies from substation to substation and from season to season, and k-means is a much more reliable algorithm. It provides only one hyperparameter ($k = 2$), it has linear computational complexity, and it has graceful behaviour if the data quality deteriorates. Although it is useful for exploratory analysis in a controlled environment, DBSCAN does not have the merit for being used in an automated, large-scale scenario.

It is worth considering the `night_mean` feature in particular due to the null result (Cohen's $d = 0.22$, $p = 0.437$) that has practical implications. The discriminative signal is in the evening window where post-commute charging takes place, and overnight baseload is not systematically affected by EV ownership. This translates to this feature extraction only needs to be done for the 18:00–24:00 window, thereby saving computational cost and data storage space for the large-scale screening across LESCO's 4.5 million connections.

5.1.3 LSTM-Based Load Forecasting Performance

On the held-out test data, the LSTM model yielded an MAE of 0.168 MVA, RMSE of 0.229 MVA, MAPE of 4.86% and R-squared of 0.974, showing improvements of 53.8% (MAE), 56.4% (RMSE)

and 52.0% (MAPE) over the Seasonal Naive baseline. The improvements are in the range of 35%–40% MAPE in comparison to LSTM-based residential forecasting reported by Kong et al. (2019) and are comparable with the general results reported by Haben et al. (2019) in their literature review, where the deep-learning-based methods consistently outperform the traditional statistical ones for short-term forecasting.

The improvement is a result of how the LSTM can capture the day-to-day load variation, which a persistence baseline cannot. The three-way chronological split (70/15/15 train/validation/test) along with an independent validation set for early stopping is used to mitigate the risk of temporal leakage in simpler two-way splits, and to ensure that the reported metrics are truly out of sample. The test period covers the hot months of Lahore before the monsoon, during which the model was presented with significant changes in the load due to weather, and the model tracked these changes with an R-squared of 0.974.

The advantage of LSTM in LESCO is due to some specific factors of the residential consumption in Pakistan. Temperature is an important exogenous factor in LHS's subtropical climate, as the temperature in summer regularly goes over 45°C and air conditioning shares in residential demand is a prominent one. The LSTM is very successful in learning the nonlinear relationship between the ambient temperature and the cooling load and turns the source of volatility into a source of predictability. The temperature sensitivity analysis revealed that the model performs well throughout the entire range of temperatures of Lahore (i.e. below 20°C and above 35°C) and did not show any significant drop in performance under extreme temperature conditions, a characteristic which Marino et al. (2016) showed was necessary for the limitative residential forecasting.

The MAE of 0.168 MVA at the transformer level corresponds to about 1.3% of the 13 MVA ONAF rating and 5.9% of the mean load, which is well below the required level of precision of Haben et al. (2019) in the context of day-ahead scheduling. This means that the 10% safety margin (the 1.3 MVA buffer between the 13 MVA rating and the 11.7 MVA effective ceiling) is unnecessarily high in this case under typical conditions, as a plus-minus 2-sigma forecast band of ~plus-minus 0.46 MVA is comfortably within the margin.

5.1.4 Grid Impact of Coordinated Charging

In addition to the two-scenario comparison (uncoordinated versus optimized) previously used in the preliminary analysis, the charging optimization evaluation also shows charging optimization value in a three-scenario comparison between Uncoordinated, FCFS, and AI-DSM Coordinated.

In the case of uncoordinated charging, 1,200 EV chargers were located during the natural evening peak, causing an aggregate load of 15.57 MVA (overload of 119.8% of ONAF rating) and an overload of 2.5 hours per day. This brought the winding hot-spot to 122.3 C - approximately 12 C above the 110C limit - and the maximum ageing-acceleration factor to 3.4 and the daily loss-of-life to approximately 28 times the no-EV baseline. This is consistent with results found by Hilshey et al. (2013) and Gong et al. (2012), which indicated that uncoordinated EV charging can lead to large reductions in transformer life, even at low levels of EV penetration.

In order to avoid any cases of overloading, the FCFS scenario was able to limit the maximum level of any individual charging session to the 13 MVA ONAF rating, in which case the loading level was kept at 100.0% or 13.0 MVA. This is a maximum decrease of 16.5% compared to the baseline without coordination. Its simplicity, only real-time monitoring and a queue, makes it an attractive interim solution; however, the transformer was loaded near its maximum, with the maximum F_{AA} of 1.03 indicating that it had intermittent ageing above the transformer's rated ageing rate, which compounds over the multi-year asset-management horizon.

The AI-DSM coordinated approach aimed to keep a safety margin and resulted in a peak reduction of 36.1% or 9.95 MVA (76.5% ONAF). More importantly, it moved EV charging to the overnight valley (00-06:00) and increased the overnight average to 5.52 MVA against FCFS's 3.44 MVA. The ageing analysis showed that the performance was almost at the baseline values ($F_{AA} = 0.073$ vs. the no-EV baseline of $F_{AA} = 0.066$), effectively maintaining the transformer's service life.

The difference between FCFS and AI-DSM highlights the importance of prediction in coordination. FCFS is a reactive approach, which means that it reacts to events that could lead to constraint violations. AI-DSM is predictive, which means that it plans to charge at times of known low demand before any violation is imminent. The reactive approach is inevitably at or near the limits of capacity during the peak period in the evening when the ambient temperatures are highest, and the thermal stress is greatest. With the predictive approach, the constrained period is avoided, and the EV load is focused during the coolest nighttime hours, where IEEE C57.91 ambient temperature corrections diminish the thermal effects for a given loading level. This time synchronization of high loading with low ambient temperature makes the AI-DSM's ageing advantage more than is implied by the load metrics.

This difference in infrastructure efficiency is captured in the load-factor metric. AI-DSM had the maximum load factor of 62.2%, while FCFS and uncoordinated charging had a load factor of 47.6% and 39.9%, respectively. This improvement from 39.9% to 62.2% is a 56% increase in efficiency of energy use – more energy delivered for the same level of installed energy, and all of this is

achieved by temporal redistribution of the same total amount of energy. This directly translates to deferred reinforcement and maximising return on LESCO's existing infrastructure, without capital investment, across LESCO's fleet of grid-station transformers.

5.2 Addressing the "Blindness" Problem

At the heart of this research was the fact that LESCO had no basic understanding of how EVs were being adopted, the “blindness” problem that was expressed in Chapter 1. This information asymmetry is also endemic to developing country distribution utilities and results in planning failures: transformer overloads seem “random” to network planners, but the network planner does not know the underlying reason for the overloads is because of EV clustering.

The identification module overcomes this blindness and achieves high identification accuracy ($F1 = 1.000$ at 30% penetration). This capability alone offers actionable intelligence, even if a complete charging coordination isn't performed. LESCO can help identify which of its 47,000 distribution transformers are serving concentrated populations of EVs and therefore can be monitored, proactively scheduled for maintenance and prioritized for capacity reinforcement. The identification module directly tackles the visibility gap of the spatial clustering of early EV adopters, as pointed out by Pearre et al. (2011), which causes localized grid stress with no system-wide metrics visible.

This shift in paradigm from “after the fact” (i.e. LESCO received reports of transformer overloads) to “before the fact” (i.e. LESCO was able to identify at-risk transformers) has changed the way the system operates. This information asymmetry was recognized by Verzijlbergh et al. (2012) as a major obstacle to successful EV integration; the identification capability shown here is possible without the need for customer self-reporting, vehicle registration databases, or new sensors and is based on the already installed smart metering infrastructure.

5.3 Implications for LESCO Implementation

The practical implications of the experimental findings are straightforward within the LESCO operational context. Based on the implementation framework suggested by Siano (2014), the following section offers a set of guidelines for implementation that are progressive over time and take into account both the technical challenges and the organizational capacity limitations of developing country utilities.

5.3.1 Phased Rollout Strategy

A staged approach is recommended, which will provide progressive development of the organization's capability together with early returns. Gellings (2009) stresses that failure is most often a problem of the sequencing of the implementation of a DSM programme rather than a technical problem.




Figure 5. 1: Proposed four-phase implementation roadmap for LESCO, transitioning from identification to full-scale deployment over a 24-month horizon.

In the initial phase (from month 1 to 6), LESCO should apply the k-means clustering algorithm on the existing Advanced Metering Infrastructure (AMI) data to create a database of households with EVs. This database should be verified with transformer serving area maps to determine equipment at an elevated stress risk. Modest computational needs—with a linear-time algorithm, 600,000 metered connections can be processed in two to three hours on commodity processors, similar to the typical monthly billing runs.

Phase 2 (Months 6-12) includes training LSTM forecasting models for the high-priority transformers identified in Phase 1 and prioritizing the 10/13 MVA grid-station power transformers for this action, as they are critical for the network. Models should be linked with the currently existing SCADA and Distribution Management Systems. Forecast accuracy should be checked using the actual measurements, and feature engineering should be adjusted based on the feedback from operations.

The third phase (months 12-18) involves choosing a pilot zone with a significant EV population to be identified and implementing the communication system necessary for charging coordination. The charging schedules as determined by the optimization module can be carried

out using smart plugs, OCPP-compliant charging stations or relay switches that are controlled by the utility. The algorithm should start off with conservative capacity limits (e.g. 15-20%, instead of 10%) to establish confidence.

Full-scale deployment (months 18-24) is the scaling-up of the framework in LESCO's service area. But for regions where the complete optimization infrastructure is not yet in place, the FCFS method is also available. Being only a real-time load monitoring tool and not a day-ahead forecasting and scheduling tool, FCFS can be implemented as an interim solution that avoids overloading without the complexity of the AI-DSM, and then further upgraded to full optimization when the infrastructure is mature.

5.3.2 Geographic Prioritization

The identification and forecasting capabilities allow the geographic prioritization of interventions, which Putrus et al. (2009) proved to be significantly more cost-effective. LESCO should first target feeders with high EV penetration (defined as $> 25\text{--}30\%$ of the feeder population) identified through clustering as having the most imminent risk. Particular attention should be given to transformers that are nearing their end of life, as even a small overload will increase the progression towards dielectric failure.

Higher EV adoption is likely to occur in upper-income neighborhoods in DHA, Gulberg, and Johar Town districts of Lahore, reflecting the demographic trends found by Plötz et al. (2014) and necessitating proactive attention. Commercial-residential mixed areas can also be quite challenging in terms of scheduling, as the commercial load during the day may limit the charging window during the night, which may necessitate new optimization formulations to include non-residential demand patterns.




Figure 5. 2: Geographic prioritization matrix guiding LESCO's DSM deployment. High EV penetration areas and critical infrastructure zones receive top priority.

5.4 Scalability Assessment

One important consideration for practical implementation is whether the framework developed on a single grid-station power transformer connected to approximately 4000 residential consumers can be scaled up to LESCO's whole service area, with 4.5 million consumers and 47 thousand distribution transformers, or to other distribution companies in Pakistan.

5.4.1 Computational Scalability

The scaling characteristics are different for the three modules. The identification module has $O(n)$ complexity per household, which can be used for batch processing. The clustering algorithm takes 2-3 hours to run on commodity hardware for LESCO's ~600,000 metered connections.

After training, the forecasting module has an $O(1)$ inference complexity per transformer. Models can be reused for transformers with similar customer profiles, and transfer learning can additionally reduce models' training needs, as did Sun et al. (2019). The optimization module is quadratic programmed per transformer but is limited by the number of EVs per transformer (usually 10-50), not by system totals. Optimization takes place locally; there is no coupling of different transformers, and the problem is inherently parallelizable.

The methodology is applicable in other DISCOs in Pakistan, since it is based upon progressively available data sources: smart meter readings, publicly available weather data and standard transformer specifications. Simple adaptations for other regions are possible; IESCO's colder winters will require reduced temperature weighting; in K-Electric's case, different baselines for the industrial, commercial and residential loads will be required; in the case of GEPCO and MEPCO, different modelling of agricultural loads will be required. These are engineering adaptations, NOT fundamental methodological limitations.

5.5 Limitations

These findings are limited by several methodological and contextual factors and need to be explicitly acknowledged.

The most critical constraint is related to data provenance. While datasets were also designed to match the consumption, data reported in LESCO's Annual Report 2022-2023 and have been carefully bottom-up modelled with appliance and grid-station aggregation, field validation with real data measured on LESCO's grid would give greater confidence to the framework's real-world performance. The identified EV charging signatures might be less pronounced in households where there are multiple high-power loads (heat pumps, electric water heaters), which are not widely present in the current Pakistani residential scenario but might gain popularity with the progress of electrification. The classification performance at 30 % penetration at 7 kW chargers might not hold at other penetration levels and other charger ratings (e.g., 3.3 kW), as the signal-to-noise ratio will significantly decrease in these cases, which is a boundary condition to be investigated systematically in future research.

The simulation assumed all users have Level 2 charging equipment with a capacity of 7 kW and have the same plug-in home arrival time patterns in the evening. In reality, charging capacities are anywhere from 3.3 kW to 22 kW (Hardman et al., 2018), and many owners may charge at their workplaces or have very different weekend/weekday charging impacts (Pearre et al., 2011). Battery degradation introduces temporal non-stationarity over the vehicle lifetime, which was not

modelled. Statistical feature engineering handling of course, some behavioral variability, but it is not yet clear how sensitive the results are to the variability of charging patterns.

A two-input LSTM network was used to train and evaluate the short-term load-forecasting model as mentioned in Section 3.5, and the result of transformer loading for the next 24 hours was predicted from the reconstructed twelve-month record. The relative performance improvement of the 53.8% MAE is important to interpret: MAE is particularly high when there is significant day-to-day weather variation in the test period (which makes the baseline look bad), and the generalisation of the model trained on a single grid-station's load history is not independently confirmed.

Optimization formulation is based on the premise that users accept utility-prescribed charging schedules. Research in behavioural economics indicates that range anxiety is sufficient to cause resistance to scheduling flexibility, even in the absence of a mobility penalty with a 100% satisfaction constraint (Axsen and Kurani, 2012). For actual deployment, incentive design research would be needed to ensure ongoing participation. The FCFS scenario (no user feedback other than plug-in) may have fewer acceptance barriers and could be a steppingstone towards full coordination.

Only modelling of the transformer thermal loading was considered as the limiting constraint, and voltage drop, protection coordination, phase imbalance and reactive power impacts were neglected. It is shown in the work of Clement-Nyns et al. (2010) and Turker et al. (2012) that the above phenomena play a more important role the higher you penetrate or the weaker the network segment. The main constraint for the modelled 30% penetration on the 10/13 MVA grid-station power transformer is thermal capacity, but it would be useful to perform a comprehensive power flow analysis for higher penetrations.

5.6 Comparison with Prior Work

The findings are generally favourable when compared with previous studies within each of the framework components, although comparisons are not easy due to a number of differences in context, data, and methodology.

The difference between the two algorithms is negligible, with both achieving an F1 of 1.000, significantly higher than the 0.85–0.92 range reported for disaggregation of European appliances by Cominola et al. (2017) and the 87–91% accuracy of the supervised disaggregation algorithms presented by Wang et al. (2018). This is due to the higher signal-to-noise ratio that results from charging at 7 kW, versus Pakistani baseline consumption of 0.7–1.4 kW, and European baselines

of 1–3 kW, where the charging signature is proportionately lower. This contextual advantage indicates that the unsupervised EV detection may be more feasible in developing countries than in developed countries, where more high-power loads are fighting for a discriminative signal in the residential environment.

The 53.8% MAE improvement is greater than the 35-40% reduction reported in Kong et al. (2019) and the improvements reported by Marino et al. (2016) for load forecasting. Due to excellent temperature-load correlation in the climate of Lahore, this temperature is an excellent exogenous predictor for the LSTM, and the R-squared value of 0.974 puts the model among the better-performing transformer-level forecasting in the literature.

The AI-DSM approach reported a peak reduction of 36.1%, which lies in the upper end of the range of 25-40% documented by Clement-Nyns et al. (2010) and is greater than the 30-35% documented by Deilami et al. (2011) for charging optimization. This is due to the significant overnight valley in Pakistani residential profiles (baseload of around 3.4 MVA against a 13 MVA transformer), which gives ample opportunity for valley filling. The three-scenario methodology (Uncoordinated / FCFS / AI-DSM) offers a detailed assessment which is more granular than the typical two-scenario comparison found in the literature.

The ageing analysis of the transformer with the IEEE C57.91 approach is something that was not primarily reported in most of the previous EV integration studies, which only provide the peak load metrics and not the long-term asset health implications. The fact that uncoordinated charging brings the winding hot-spot up to almost 122 °C, some 12 °C over the rated limit of 110 °C, and increases the daily loss-of-life nearly 28-fold when compared to the no-EV baseline, whereas the AI-DSM keeps ageing near to the no-EV baseline (maximum $F_{AA} \approx 0.07$), gives an economic incentive for coordination, as well as an operational convenience.

5.7 Summary

This chapter has presented results from Chapter 4 within the context of the literature, practical considerations for deployment, and identified limitations. The AI-DSM framework illustrates the potential of intelligent, software-enabled DSM to turn the expected problem of EV integration into an opportunity for a new, manageable phase in the evolution of LESCO's distribution network. The identification module addresses the blindness issue with regard to proactive planning. The forecasting module ensures the forecasting intelligence for optimum scheduling. The optimization module enables

intelligence to be translated into actionable charging schedules to protect infrastructure, optimize asset utilization and ensure user satisfaction.

The three-tier evaluation (Uncoordinated → FCFS → AI-DSM) shows that simple reactive coordination removes the acute problem of transformer overloading, while predictive optimization ensures the full span of benefits—safety margins against uncertainty, near baseline transformer ageing, and maximum infrastructure utilization. The framework is an approach that provides a capital-efficient path to enabling LESCO and other utilities in a developing country to transition to transport electrification without the need for large-scale infrastructure investments, relying on data and algorithms instead of steel and copper.

References

- [1] 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.
- [2] 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.
- [3] 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 the voltage profile. *IEEE Transactions on Smart Grid*, 2(3), 456-467.
- [4] Estebarsari, A., Barbato, A., & Pons, E. (2014). Appliance classification via a clustering-based approach. 2014 IEEE International Conference on Smart Grid Communications (SmartGridComm), 326-331.
- [5] 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.
- [6] Grey, M. K., & Morsi, W. G. (2015). Power quality assessment in distribution systems impacted by plug-in hybrid electric vehicles. *IEEE Transactions on Smart Grid*, 6(2), 660-668.
- [7] Hart, G. W. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12), 1870-1891.
- [8] Hilshey, A. D., Hines, P. D., Rezaei, P., & Dowds, J. R. (2013). Estimating the impact of electric vehicle smart charging on distribution transformer ageing. *IEEE Transactions on Smart Grid*, 4(2), 905-913. <https://doi.org/10.1109/TSG.2012.2217385>
- [9] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
- [10] IEEE Standards Association. (2011). IEEE Standard C57.91-2011: Guide for loading mineral-oil-immersed transformers and step-voltage regulators. IEEE.
- [11] International Energy Agency. (2023). Global EV Outlook 2023: Catching up with climate ambitions. IEA Publications.
- [12] Kong, W., Dong, Z. Y., Jia, Y., Hill, D. J., Xu, Y., & Zhang, Y. (2019). Short-term residential load forecasting based on an LSTM recurrent neural network. *IEEE Transactions on Smart Grid*, 10(1), 841-851. <https://doi.org/10.1109/TSG.2017.2753802>

- [13] Lahore Electric Supply Company. (2023). Annual Report 2022-2023. LESCO Publications.
- [14] 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.
- [15] Ministry of Climate Change. (2019). National Electric Vehicle Policy 2019. Government of Pakistan.
- [16] Muratori, M. (2018). Impact of uncoordinated plug-in electric vehicle charging on residential power demand. *Nature Energy*, 3(3), 193-201.
- [17] Parson, O., Ghosh, S., Weal, M., & Rogers, A. (2012). Non-intrusive load monitoring using prior models of general appliance types. *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1).
- [18] Raza, W. Q., & Khosravi, A. (2015). A review of artificial intelligence-based load demand forecasting techniques for smart grid and buildings. *Renewable and Sustainable Energy Reviews*, 50, 1352-1372.
- [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] 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. <https://doi.org/10.1109/TSG.2013.2251483>
- [21] Wang, Y., Chen, Q., Gan, D., Yang, J., Kirschen, D. S., & Kang, C. (2019). Deep learning-based socio-demographic information identification from smart meter data. *IEEE Transactions on Smart Grid*, 10(3), 2593-2602. <https://doi.org/10.1109/TSG.2018.2805723>
- [22] 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.
- [23] Zheng, J., Xu, C., Zhang, Z., & Li, X. (2017). Electric load forecasting in smart grids using long-short-term-memory based recurrent neural network. 2017 51st Annual Conference on Information Sciences and Systems (CISS), 1-6.
- [24] Ahmad, T., Chen, H., Guo, Y., & Wang, J. (2018). A comprehensive overview of the data-driven and large-scale approaches for forecasting of building energy demand: A review. *Energy and Buildings*, 165, 301-320.

- [25] Ahmad, W., Ayub, N., Ali, T., Irfan, M., Awais, M., Shiraz, M., & Glowacz, A. (2020). Towards short-term electricity load forecasting using improved support vector machine and extreme learning machine. *Energies*, 13(11), 2826.
- [26] Axsen, J., & Kurani, K. S. (2012). Interpersonal influence within car buyers' social networks: Applying five perspectives to plug-in hybrid vehicle drivers. *Environment and Planning A*, 44(5), 1047-1065.
- [27] 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.
- [28] Gellings, C. W. (2009). *The Smart Grid: Enabling Energy Efficiency and Demand Response*. CRC Press.
- [29] 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.
- [30] Hardman, S., Jenn, A., Tal, G., Axsen, J., Beard, G., Daina, N., ... & Witkamp, B. (2018). A review of consumer preferences for and interactions with electric vehicle charging infrastructure. *Transportation Research Part D: Transport and Environment*, 62, 508-523.
- [31] Hidrue, M. K., Parsons, G. R., Kempton, W., & Gardner, M. P. (2011). Willingness to pay for electric vehicles and their attributes. *Resource and Energy Economics*, 33(3), 686-705.
- [32] Hu, J., Morais, H., Sousa, T., & Lind, M. (2016). Electric vehicle fleet management in smart grids: A review of services, optimization and control aspects. *Renewable and Sustainable Energy Reviews*, 56, 1207-1226.
- [33] Kempton, W., & Tomić, J. (2005). Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1), 268-279.
- [34] Khattak, S. H., Arentsen, M., & Rozema, J. (2016). Electricity distribution company performance in developing countries: A case study of LESCO. *Energy Policy*, 96, 485-493.
- [35] Lusi, P., Khalilpour, K. R., Andrew, L., & Liebman, A. (2017). Short-term residential load forecasting: Impact of calendar effects and forecast granularity. *Applied Energy*, 205, 654-669.
- [26] Muratori, M., & Rizzoni, G. (2016). Residential demand response: Dynamic energy management and time-varying electricity pricing. *IEEE Transactions on Power Systems*, 31(2), 1108-1117.

- Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3), 381-388.
- Pearre, N. S., Kempton, W., Guensler, R. L., & Elango, V. V. (2011). Electric vehicles: How much range is required for a day's driving? *Transportation Research Part C: Emerging Technologies*, 19(6), 1171-1184.
- Plötz, P., Schneider, U., Globisch, J., & Dütschke, E. (2014). Who will buy electric vehicles? Identifying early adopters in Germany. *Transportation Research Part A: Policy and Practice*, 67, 96-109.
- Putrus, G. A., Suwanapingkarl, P., Johnston, D., Bentley, E. C., & Naber, R. K. (2009). Impact of electric vehicles on power distribution networks. 2009 IEEE Vehicle Power and Propulsion Conference, 827-831.
- Siano, P. (2014). Demand response and smart grids—A survey. *Renewable and Sustainable Energy Reviews*, 30, 461-478.
- Sortomme, E., & El-Sharkawi, M. A. (2012). Optimal scheduling of vehicle-to-grid energy and ancillary services. *IEEE Transactions on Smart Grid*, 3(1), 351-359.
- Sun, M., Konstantelos, I., & Strbac, G. (2019). A deep learning-based feature extraction framework for system security assessment. *IEEE Transactions on Smart Grid*, 10(5), 5007-5020.
- Turker, H., Bacha, S., Chatroux, D., & Hably, A. (2012). Low-voltage transformer loss-of-life assessment based on peak load profile analysis. *IEEE Transactions on Power Delivery*, 27(3), 1347-1355.
- Verzijlbergh, R. A., Grond, M. O. W., Lukszo, Z., Slootweg, J. G., & Ilic, M. D. (2012). Network impacts and cost savings of controlled EV charging. *IEEE Transactions on Smart Grid*, 3(3), 1203-1212.
- Wang, Y., Chen, Q., Hong, T., & Kang, C. (2018). Review of smart meter data analytics: Applications, methodologies, and challenges. *IEEE Transactions on Smart Grid*, 10(3), 3125-3148.
- Zoha, A., Gluhak, A., Imran, M. A., & Rajasegarar, S. (2012). Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors*, 12(12), 16838-16866.