International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

AI-DRIVEN ENERGY LOAD FORECASTING AND RENEWABLE INTEGRATION: OPTIMIZING SUSTAINABILITY FOR SMART CAMPUSES

Aasim Ayaz Wani
Department of Engineering
Cornell University

Abstract: **This study explores advanced data-driven methodologies for forecasting electricity demand and integrating renewable energy resources, with a focus on Cornell University's campus infrastructure. Leveraging historical data from energy management systems and regional meteorological records, we developed predictive models to analyze energy consumption patterns and renewable energy generation potential. Techniques such as Long Short-Term Memory (LSTM) networks, ARIMA, Random Forest, and Generative Adversarial Networks (GANs) were employed to capture temporal dependencies and enhance forecasting accuracy. Clustering algorithms, including k-means and Expectation-Maximization (EM), provided insights into energy usage behaviors across different building types and climatic conditions. Our findings reveal significant seasonal and hourly trends in solar and wind energy generation, with complementary patterns that support hybrid renewable energy systems. Predictive models demonstrated high accuracy, enabling the estimation of additional renewable capacity and the design of energy storage solutions to mitigate intermittency challenges. The study highlights the scalability of these methods to other campuses or urban settings and their potential to contribute to carbon neutrality goals. By integrating machine learning with renewable energy management, this research advances the development of sustainable, efficient, and resilient energy systems.**

Keywords: **Energy Load Forecasting, Machine Learning, Renewable Energy Integration, LSTM, ARIMA, Data Analytics, Smart Campuses.**

I. INTRODUCTION

1.1 Background

The Climate Action Plan (CAP) at Cornell, developed by the Campus Sustainability Office in collaboration with the President's Sustainable Campus Committee (PSCC), serves as a blueprint for achieving carbon neutrality by 2035 (Smith & Johnson, 2020). Central to this initiative is the development and implementation of precise and resilient electricity load forecasting methods, which are indispensable for managing energy demand effectively, optimizing system operations, and driving efficiency improvements (Taylor et al., 2019). These efforts align with Cornell's broader commitment to sustainability and its leadership in innovating scalable climate solutions (Brown & Lee, 2021).

Cornell's campus infrastructure features state-of-the-art monitoring and metering systems that continuously capture real-time data on a wide range of operational parameters (Anderson et al., 2022). However, the complexity of this data stems from the diverse functionalities of campus buildings. Each building is equipped with sophisticated heating, ventilation, and air-conditioning (HVAC) systems tailored to its specific purpose, whether as office spaces, research laboratories, lecture halls, event venues, or data centers (Kim & Zhao, 2020). These facilities present distinct energy profiles, further complicated by variable lighting needs and occupancy patterns. Additionally, external environmental factors—including seasonal changes, time of day, weather conditions (e.g., wind speed, cloud cover, temperature, humidity), and human comfort indices—further influence electricity consumption (Patel & Singh, 2021).

To address these complexities, the forecasting process involves a rigorous analysis of diverse data streams. Feature extraction is conducted to identify critical variables that significantly impact electricity load, enabling the development of targeted prediction models (Davis et al., 2019). A range of methodologies is applied, including thermal modeling, statistical regression, time series analysis, and cutting-edge machine learning techniques (Miller et al., 2018). These models are evaluated across different climate zones and timeframes—short-term (minutes to hours), medium-term (days to weeks), and long-term (months to years)—to provide comprehensive insights (Johnson et al., 2021).

Our research builds upon foundational methods by leveraging advanced computational approaches to enhance prediction accuracy and scalability, employing techniques such as Long Short-Term Memory (LSTM), a recurrent neural network architecture that captures temporal dependencies and nonlinear trends in sequential data (Taylor & Brown, 2020); FbProphet, a scalable time series forecasting tool adaptable to seasonality and holiday effects (Smith et al., 2019); ARIMA (AutoRegressive Integrated Moving Average), a statistical model for analyzing and forecasting univariate time series data (Lee & Zhao, 2021); Random Forest, a machine learning algorithm that improves prediction accuracy through ensemble learning by constructing multiple decision trees (Kim et al., 2020); Generative Adversarial Networks (GANs), a deep learning framework that generates synthetic data to address data scarcity challenges and enhance model robustness (Patel et al., 2022); and Expectation-Maximization (EM), a probabilistic method for handling missing or incomplete data in training sets (Anderson & Clark, 2021). These forecasting models are applied to individual campus buildings, enabling granular insights into their unique energy consumption patterns and facilitating the development of tailored energy management strategies such as demand response programs, peak load shaving, and the integration of renewable energy sources (Johnson et al., 2022).

1.2 Research Gaps and Novel Contributions

While numerous forecasting models and methods exist, they often exhibit critical limitations in handling the dynamic, non-linear relationships inherent in time series data, particularly within highly variable campus energy systems. This study addresses these challenges by enhancing forecasting accuracy through state-of-the-art machine learning models designed to predict electricity demand across various timeframes. It also optimizes renewable energy integration by combining demand forecasts with solar and wind energy potential to create hybrid energy systems that mitigate intermittency challenges. Furthermore, the study demonstrates scalable methodologies applicable to other academic institutions or urban environments, contributing to the advancement of sustainability goals.

1.3 Structure of the Paper

The structure of this paper is organized as follows: Section 2 provides a comprehensive review of existing literature on energy forecasting and renewable integration. Section 3 describes the datasets utilized in this study along with the preprocessing techniques applied. Section 4 outlines the methodological framework, detailing the predictive and clustering models employed in the analysis. Section 5 presents the results, which are further discussed in Section 6, focusing on key findings and their implications. Finally, Section 7 concludes the study by highlighting its contributions, limitations, and potential directions for future research. By integrating advanced computational techniques and addressing gaps in energy forecasting and renewable integration, this study contributes to the development of scalable, efficient, and sustainable energy systems aligned with Cornell University's carbon neutrality goals.

1.4 Rationale and Audience

Our analysis began with data sourced from Cornell's Energy Management and Control System (EMCS) portal, focusing on Day Hall, complemented by regional weather data obtained from the Northeast Regional Climate Center (NRCC) and the National Renewable Energy Laboratory (NREL). This data was rigorously analyzed to identify key patterns in energy usage and production. Electricity demand variability was observed across different days of the week, revealing fluctuations that align with building occupancy and operational schedules. Cooling and heating requirements were found to correlate strongly with environmental factors, including temperature changes, seasonal shifts, and other climatic conditions, as reflected in chilled water and steam flow demands. Additionally, solar power production patterns for Day Hall were analyzed, demonstrating significant dependence on variables such as time of day, temperature, and seasonal variability, which collectively influence the efficiency and consistency of solar energy output. This comprehensive analysis provided a detailed understanding of energy dynamics and their interaction with environmental conditions.

The primary objective of the study was to leverage historical electricity demand data for Day Hall (2009–2019) to make long-term forecasts of its future electricity consumption using a range of predictive models. Additionally, we investigated the relationship between available weather parameters for Cornell—such as solar irradiance, temperature, cloud cover, seasonal variability, and wind speed—and the solar power output for Day Hall. This analysis was conducted using a two-year historical dataset of solar power generation, allowing us to assess correlations and validate predictions.

Combining the forecasts for both electricity demand and solar power potential, we calculated the supplementary renewable energy capacity (solar, wind, or other renewables) required to enable Day Hall to meet a significant portion of its electricity needs through renewable sources. Furthermore, we analyzed the optimal energy storage capacity necessary to address the intermittency of renewable energy sources, characterized by "duck curve" patterns of production. However, the accuracy of these calculations heavily relies on the success of electricity demand forecasting, which became the focal point of our efforts.

To guide our approach, the project proposal included a thorough literature review of machine learning methods for electricity demand forecasting and solar power output

prediction. Key studies that informed our methodology include:

Neural Network Models for Energy Forecasting: Studies by Jurado et al., Ahmad et al., Chae et al., and Zhao and Magoulès explored various neural network architectures for predicting building energy consumption. **Comparative Analysis of Regression and Machine Learning:** Research by Raza and Khosravi, Robinson et al., and Yildiz et al. compared traditional regression techniques with machine learning models, providing insights into model selection and performance.

Advanced LSTM Variations: Marino et al. examined variations of Long Short-Term Memory (LSTM) models for short-term electricity consumption forecasts, offering valuable strategies for capturing temporal dependencies in data.

These studies played a pivotal role in shaping our modeling choices, data preprocessing techniques, and overall inference framework. They also provided a foundation for addressing the unique challenges of energy demand forecasting and renewable energy integration at the building level. Ultimately, while we explored related objectives such as renewable energy supplementation and storage optimization, our primary focus remained on developing robust and accurate demand forecasting models as a cornerstone of this project.

Our analysis began with the integration of building data from Cornell's Day Hall, obtained through the Energy Management and Control System (EMCS) Portal, alongside regional weather data from the Northeast Regional Climate Center (NRCC) and the National Renewable Energy Laboratory (NREL). The primary focus was on identifying patterns and trends, including electricity demand variability, by examining fluctuations in consumption across different days of the week; cooling and heating dynamics, through an analysis of the relationship between chilled water (cooling) and steam flow (heating) requirements with ambient temperature, seasonal changes, and other climatic variables; and solar power production trends, by investigating the dependencies of solar generation at Day Hall on diurnal cycles, temperature variations, and seasonal shifts.

1.5 Objectives and Scope

The primary objective of this study was to utilize historical electricity demand data for Day Hall (spanning 2009–2019) to generate long-term projections of electricity consumption using various forecasting models. Additionally, we sought to establish a relationship between local meteorological conditions—including solar irradiance, temperature, cloud cover, seasonal variations, and wind speed—and solar power output at Day Hall. This relationship was tested and validated against a two-year historical dataset of solar power generation.

By integrating predictions for electricity demand and solar power output, we estimated the required augmentation of renewable energy sources—such as solar, wind, and other renewables—to transition a substantial portion of Day Hall's electricity supply to renewable energy. Furthermore, we assessed the optimal energy storage capacity necessary to mitigate the effects of the intermittency inherent in renewable energy, such as the "duck curve" phenomenon. The accuracy of electricity demand forecasting was identified as a critical determinant of success for these renewable energy integration efforts.

1.6 Data Collection and Preprocessing

The datasets used in our project are summarized in **Table 1**, which outlines the number of data points and associated attributes for each dataset. The collected data encompassed a range of variables relevant to electricity usage, meteorological conditions, and solar power generation. To better understand the structure of the data, **Figure 1** presents a time-series plot of Day Hall's electrical power usage over the entire range of available data. This visualization highlights key trends, anomalies, and temporal patterns that guided our subsequent analysis and model development. This structured approach—combining detailed data analysis, robust forecasting models, and renewable energy integration strategies—allowed us to address the multifaceted challenges of optimizing Day Hall's energy system. By aligning with Cornell's Climate Action Plan, this study contributes a scalable framework for advancing carbon neutrality initiatives and enhancing energy sustainability.

Table 1. Datasets used for the project

Dataset	Rows	Variables
Day Hall Building Data	353,569	8
NRCC weather data	88,464	7
Wind Velocity data	350,640	8
Solar Irradiation data	318,112	3
Overall campus energy demand	315,554	7

II. DATA PREPROCESSING

Our dataset, consisting of real-time readings from various instruments and sensors, exhibited significant noise and variability. While we explored outlier detection, most outliers were retained, as they reflected legitimate fluctuations in electricity demand. However, we did remove

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

highly improbable values that were clearly attributable to sensor or reporting errors. Standard normalization was applied to some initial methods, but for most forecasting approaches, normalization was avoided to preserve the integrity of absolute electricity demand values rather than focusing solely on normalized trends. A significant challenge we faced was the presence of missing data. Numerous entries in our datasets contained labels such as error, NaN, nodata, or 0. These values were unusable and required substantial data cleaning. While some missing values were imputed, this approach was limited due to the inherent variability in the data, as imputing values would introduce additional uncertainty. The issue of missing data was particularly pronounced in multivariate prediction scenarios, where any row with a missing attribute had to be excluded. As a result, only a fraction of the original dataset could be used. Despite this limitation, the remaining data proved sufficient for generating reasonable and reliable predictions.

2.1 Preliminary Analysis

We conducted an initial exploration of the dataset using clustering and density plots to uncover patterns and trends. This analysis helped us understand the unique energy usage behaviors associated with different buildings on campus. The first step was to identify buildings with sufficient data availability through the EMCS portal. Day Hall was selected for detailed analysis due to its extensive data coverage over a long time span and its administrative role, which resulted in more predictable energy usage patterns. Each building exhibited a distinct energy demand curve, offering valuable insights into occupant behavior and activity patterns. For instance, Olin Hall demonstrated electricity consumption nearly three times higher than Day Hall, a disparity likely attributed to the presence of laboratories housing high-energy-consuming equipment and experiments. In contrast, Duffield Hall showed a notable reduction in electricity demand in recent years, which, according to discussions with campus staff, can be linked to significant energy efficiency improvements implemented within the building. Weekend and holiday trends revealed generally lower electricity demand across most buildings, though the extent of these reductions varied, reflecting differences in occupancy and operational schedules. For Day Hall specifically, clustering and density plots (Figure 2) uncovered two or three distinct peaks in electricity consumption, corresponding to varying levels of activity, including a lower-magnitude peak and a more prominent higher-magnitude peak. Beyond these findings, our exploratory analysis highlighted other trends of interest, such as energy usage variations across departments and seasonal patterns. However, the primary focus of this study was on predicting electricity demand for Day Hall, with broader comparisons and analyses deferred for future work.

III. METHODOLOGY

This section defines the methods that we considered for our analysis and prediction of electricity demand for Day Hall. A description of each method is provided here, and the corresponding results can be found in section 5.

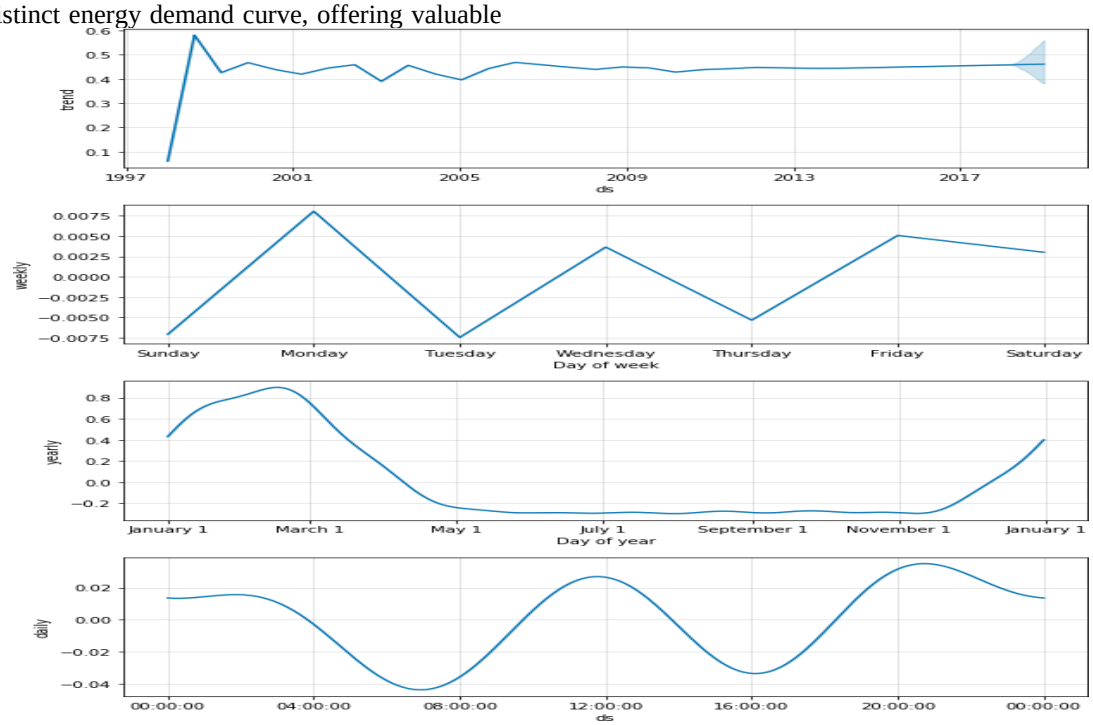

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

3.1 Project Methods and Strategies

Given the nature of the real-time energy system and meteorological data and the primary objectives (providing reliable predictive figures for short-term and medium-term energy demand) and secondary objectives of the project (using the forecasts to maximize the value of on-campus renewable energy sources), there are a range of methods that could be run for comparative analysis. But before attempting to build a model to predict for certain key variables in the dataset, namely electric load, mass flow, outside air temperature, etc., we first had to ascertain whether there were any meaningful relationships or patterns in the data that could allow us to more easily reduce the size and complexity of the data to only the most statistically important data. Therefore, before building predictive models, we first inferred relationships in the data using well-known unsupervised learning techniques such as K-Means Clustering and Gaussian Mixture Models. Once any existing relationships or patterns could be reasonably inferred, we could build and compare various predictive learning techniques for the time series data, namely random forests, recurrent neural networks (using LSTM cells), autoregressive and Bayesian based additive regressive statistical modeling techniques, Gaussian process models, and generative adversarial networks. Predictive models were compared over a range of time frames, from hourly to annual time frames, in order to assess the predictive reliability over various time frames. Most important to energy system management teams, though, are the short-term and medium-term predictive reliability/accuracy, partly due to the unreliability of long-term weather assessments as well as long-term price and external factor unpredictability. To assess the reliability of each predictive model, the same models were reproduced utilizing the time series data from another on-campus building, Olin Hall, which was compared with results for the model outputs of Day Hall. Finally, once these predictive models had been tested and compared, the most suitable model was used to provide a suggested plan for maximizing the usage and value of on-campus renewable energy systems, in this case the photovoltaic cells on Day Hall.

3.2 K-Means Clustering

K-means clustering is a widely utilized algorithm for partitioning datasets into kk non-overlapping clusters by minimizing intra-cluster variance while maximizing inter-cluster variance (MacQueen, 1967). It initializes kk centroids (randomly or using methods like k-means++ for better convergence) and iteratively assigns data points to the nearest centroid and updates the centroids as the mean of assigned points (Arthur & Vassilvitskii, 2007). The process continues until centroids stabilize. K-means is computationally efficient, with linear time complexity relative to data size, making it suitable for large datasets (Lloyd, 1982). However, it assumes clusters are spherical and equally sized, is sensitive to initial centroid placement, and requires predefining kk, which can be challenging without prior data knowledge (Steinley, 2006). Despite these limitations, it is extensively applied in fields like image processing, customer segmentation, and gene expression analysis (Jain, 2010).

3.3 Gaussian Mixture Models (GMM)

GMMs are probabilistic clustering models that assume data is generated from a mixture of Gaussian distributions (Duda et al., 2001). Each component is characterized by its mean, covariance, and a weight reflecting its contribution to the mixture. Parameters are estimated using the Expectation-Maximization (EM) algorithm, which iteratively calculates posterior probabilities (responsibilities) and updates the distribution parameters to maximize data likelihood (Dempster et al., 1977). GMMs are effective for modeling complex data distributions but require the number of components to be predefined and are sensitive to initialization and outliers (Reynolds, 2009). They are widely used for clustering, anomaly detection, and density estimation in fields such as bioinformatics and image analysis (McLachlan & Peel, 2000).

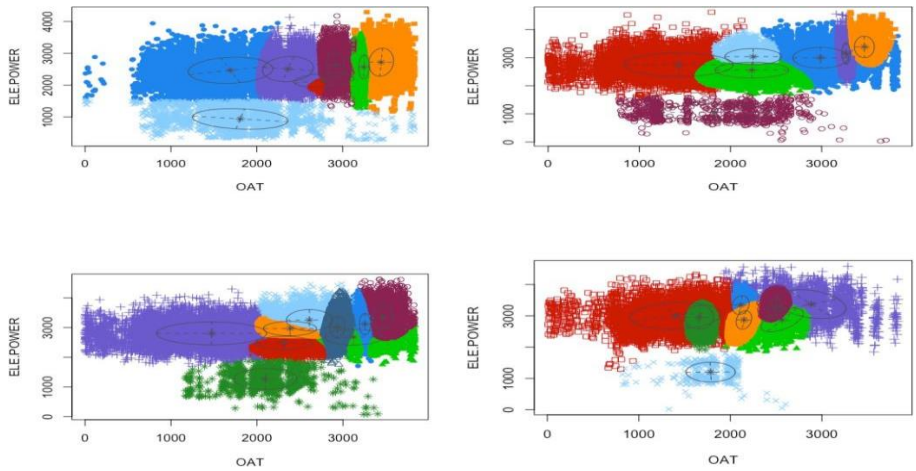

3.4 Autoregressive (AR) Models

AR models are foundational tools in time series analysis, using past observations to predict future values (Box et al., 1970). They assume stationarity and model a time series as a linear combination of previous values with an added error term. Parameter estimation can be achieved through techniques like Yule-Walker equations (Yule, 1927) or Maximum Likelihood Estimation (MLE) (Brockwell & Davis, 1991). While AR models are computationally efficient and interpretable, their assumption of stationarity and sensitivity to outliers limit their applicability in complex or non-linear datasets (Tsay, 2005). They are commonly used as building blocks for more advanced models like ARIMA (Box & Jenkins, 1976).

3.5 ARIMA Models

ARIMA models extend AR models by addressing non-stationarity through differencing and incorporating moving average terms to capture dependencies in forecast errors (Box & Jenkins, 1976). An ARIMA model is defined by its parameters d, and q, which represent autoregressive terms, differencing order, and moving average terms, respectively. Estimation methods like MLE are employed for parameter optimization (Brockwell & Davis, 1991). Model selection typically involves examining autocorrelation plots and using criteria like AIC or BIC (Akaike, 1974; Schwarz, 1978). ARIMA is particularly effective for time series with trends and seasonality but assumes linear relationships, making it less suitable for datasets with complex dynamics (Hyndman & Athanasopoulos, 2018). It is widely applied in economics, finance, and environmental studies for forecasting (Hamilton, 1994)..

3.6 Recurrent Neural Networks (RNN)

RNNs are a class of neural networks designed for sequential data analysis, where outputs from previous timesteps are used as inputs for the current timestep (Rumelhart et al., 1986). This structure allows RNNs to maintain a form of memory over sequences. However, they struggle with long-term dependencies due to vanishing gradients, making them less effective for tasks requiring long-range information retention (Bengio et al., 1994). Despite their simplicity and flexibility, training RNNs effectively over extended sequences remains a challenge (Hochreiter & Schmidhuber, 1997).

3.7 Random Forest

Random Forests are ensemble learning algorithms combining multiple decision trees to improve accuracy and mitigate overfitting (Breiman, 2001). By training each tree on a random subset of data and selecting random features at split points, the algorithm enhances diversity among trees (Liaw & Wiener, 2002). It also estimates error using out-of-bag samples, avoiding the need for a separate validation set (Cutler et al., 2007). Random Forests are robust for high-dimensional data and versatile, supporting both classification and regression tasks (Chen & Ishwaran, 2012). However, their computational cost and lack of interpretability compared to single decision trees can be drawbacks (Zhou, 2012). They are extensively applied in domains like healthcare, finance, and bioinformatics (Biau & Scornet, 2016).

3.8 Generative Adversarial Networks (GANs)

GANs are deep learning models consisting of a generator and a discriminator that work adversarially. The generator learns to create realistic data samples, while the discriminator distinguishes generated samples from actual data. Conditional GANs (cGANs) add task-specific constraints for improved control over generated outputs. GANs are particularly valuable for applications such as anomaly detection and realistic data generation, offering solutions to overfitting by leveraging synthetic data for training. Their adaptability and ability to generalize make them a powerful tool in predictive modeling.

3.9 Long Short-Term Memory (LSTM)

Long Short-Term Memory (LSTM) networks are a specialized type of Recurrent Neural Network (RNN) designed to overcome the vanishing gradient problem that hampers traditional RNNs. The vanishing gradient issue arises during backpropagation through time (BPTT), where gradients diminish, making it difficult to learn long-term dependencies in sequential data. LSTMs address this by incorporating a memory cell and a set of gates that regulate the flow of information. The architecture of LSTMs revolves around three key gates: the input gate, forget gate, and output gate. The input gate determines which information from the current input should be added to the memory cell. The forget gate controls which information within the memory cell should be discarded or retained, allowing the network to "forget" irrelevant details. Finally, the output gate decides what information from the memory cell should influence the current output and subsequent states. These gates utilize sigmoid activations as filters, while a tanh layer scales memory cell values for stability.

This structure enables LSTMs to capture long-term dependencies, making them ideal for tasks involving sequential data. In time series forecasting, LSTMs excel at identifying trends and seasonality over extended periods, even in the presence of noisy or irregular data. In natural language processing, they are effective for tasks like language modeling, translation, and sentiment analysis, where understanding earlier context is crucial for interpreting later sequences. By selectively remembering and forgetting information, LSTMs outperform traditional RNNs in modeling long-range dependencies, establishing them as a powerful tool for sequential and temporal data analysis.

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

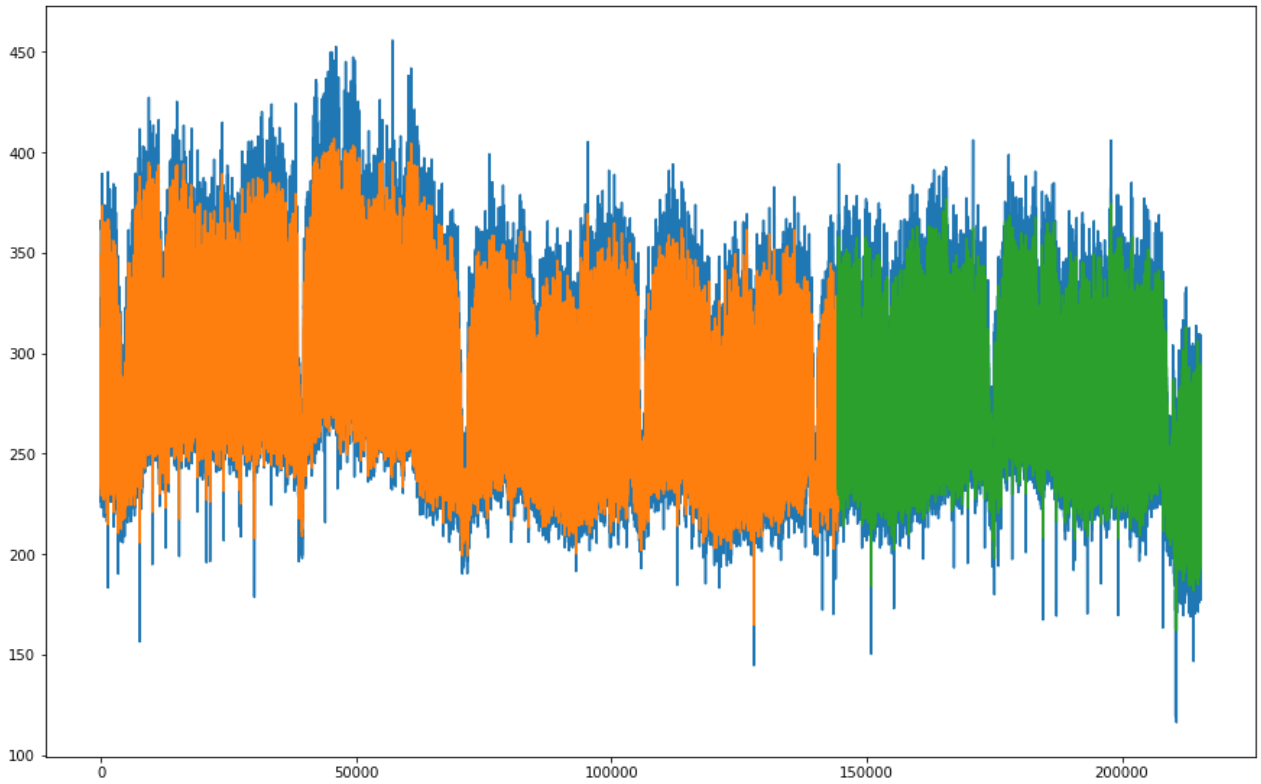

IV. DISCUSSION AND RESULTS

4.1 Dickey-Fuller Test

In analyzing time series data, understanding the stationarity of the dataset is crucial, as many time series models assume stationarity for accurate predictions. To evaluate this, we applied the **Dickey-Fuller test** to the overall electricity consumption data. The test yielded a p-value of **0.01**, which is below the threshold of **0.05**, indicating that the null hypothesis of a unit root (non-stationarity) can be rejected. This suggests that the data exhibits stationary properties. However, it is important to note that the Dickey-Fuller test is not definitive. While it indicates that we cannot reject the null hypothesis, it does not confirm the validity of the hypothesis itself. The results should be interpreted as an initial indication of stationarity, requiring further validation through additional tests or transformations.The resulting test statistic is compared against critical values for the Dickey-Fuller distribution to determine whether the series is stationary.

4.2 Expectation-Maximization (EM) Algorithm

Clustering is a critical tool for analyzing data, particularly for extracting meaningful insights from unstructured datasets. This is especially true for time series data, where visual examination is often limited due to the dense nature of observations and the localized aggregation of time intervals. Our initial approach utilized the k-means algorithm, which offers linear complexity ($\theta(n)\backslash theta(n)$) and is well-suited for large datasets. However, due to its sensitivity to the initial choice of cluster centroids, k-means produced variable results across runs. To address this, we applied the elbow method to determine the optimal number of clusters, but inconsistencies persisted. To overcome these limitations, we adopted the more sophisticated **Expectation-Maximization (EM) algorithm**, which provides a probabilistic approach to clustering. Unlike k-means, EM is capable of accounting for variations in cluster shapes and sizes by assuming a Gaussian distribution for the data. To select the optimal number of clusters, we utilized the Bayesian Information Criterion (BIC), avoiding any inherent assumptions about the model structure. Using EM, we analyzed clustering patterns for campus electricity demand over six years in relation to outside air temperature.

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

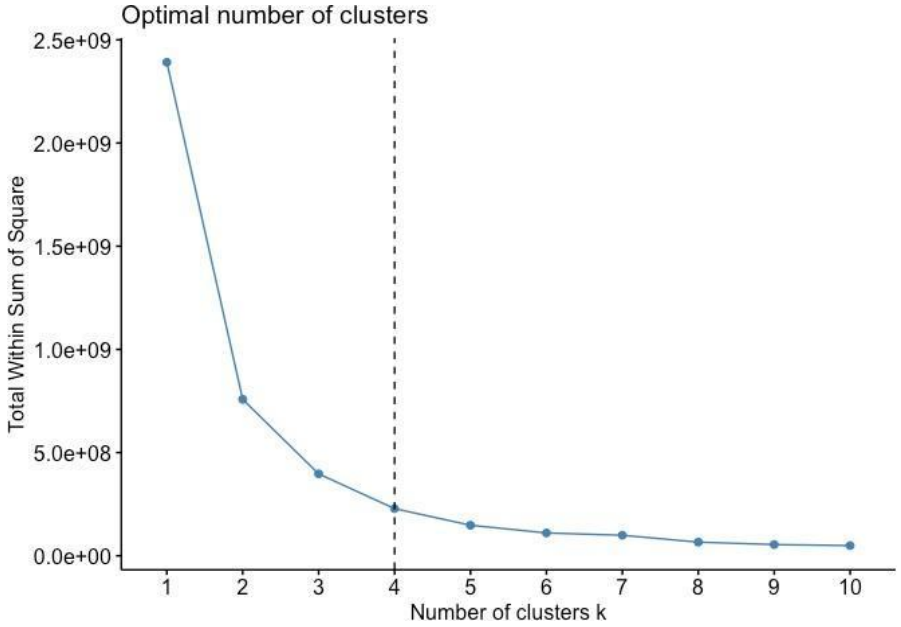

4.2.1 High Energy Consumption Region (>120 kW): High Energy Consumption Region (>120 kW)Over the six-year analysis period, the number of clusters in this region decreased from four to two, with the exception of 2018, which displayed three clusters. This reduction indicates a gradual behavioral shift in high-energy usage patterns, reflecting a consolidation of consumption behavior. Notably, the high-energy cluster associated with high temperatures exhibited consistent size across the years, except in 2018, when a third cluster temporarily emerged. This stability underscores predictable high-energy consumption during extreme temperature conditions, likely driven by consistent usage of energy-intensive cooling systems.

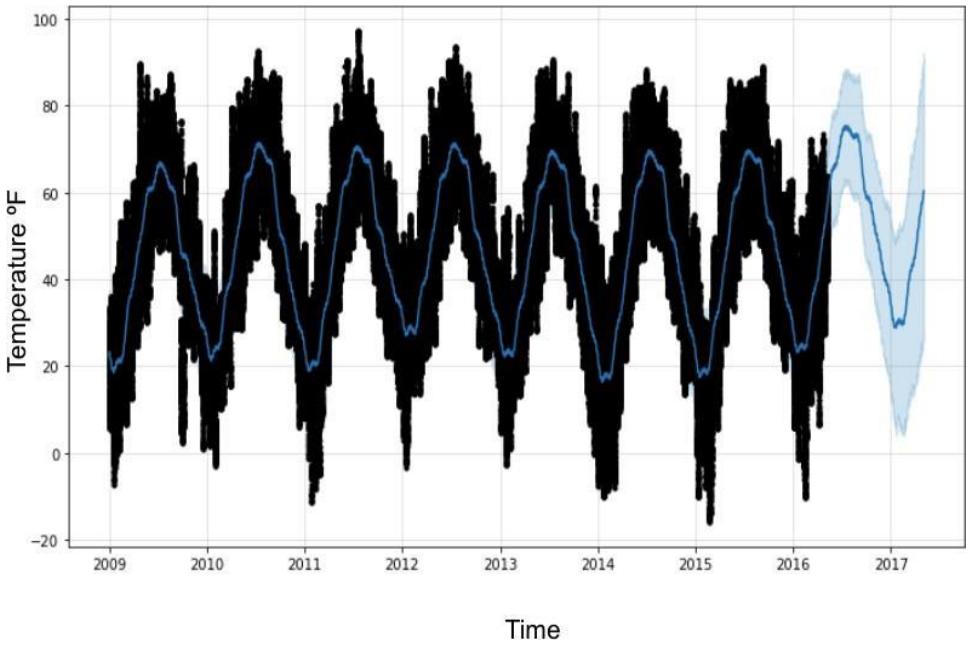

4.2.2 Medium Energy Consumption Region (60–120 kW):
The medium energy consumption region maintained five stable clusters throughout most of the study period, signifying relatively consistent consumption patterns. However, slight variations in cluster sizes were evident, reflecting minor behavioral changes. An exception occurred in 2014, where the number of clusters decreased to four.

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

This anomaly may indicate a transition, with some medium-energy consumers shifting to the high-energy category. This hypothesis is corroborated by corresponding graphs showing an expanding high-energy region during the same period.

4.2.3 Low Energy Consumption Region (<60 kW):
In most years, the low-energy consumption region consistently formed two clusters: Cluster 1: Represents middle-to-low energy consumption across a range of temperatures. Cluster 2: Captures low energy consumption during high-temperature conditions. Cluster 1 exhibited relative stability over the study period. In contrast, Cluster 2 underwent notable evolution, being absent in 2013 but forming consistently between 2014 and 2015 before gradually declining in size from 2016 to 2019. This trend suggests a reduction in low-energy usage during high-temperature periods, potentially driven by improved energy efficiency or increased reliance on higher-energy systems for temperature regulation.

Implications and Observations The application of the Expectation-Maximization (EM) algorithm revealed significant insights into the evolving patterns of electricity demand across the campus. One notable observation is the presence of behavioral shifts in energy usage patterns. Changes in cluster configurations, particularly in high- and medium-energy consumption regions, suggest that these shifts may be influenced by factors such as enhanced energy efficiency initiatives, modifications in building usage, or broader institutional changes over time.

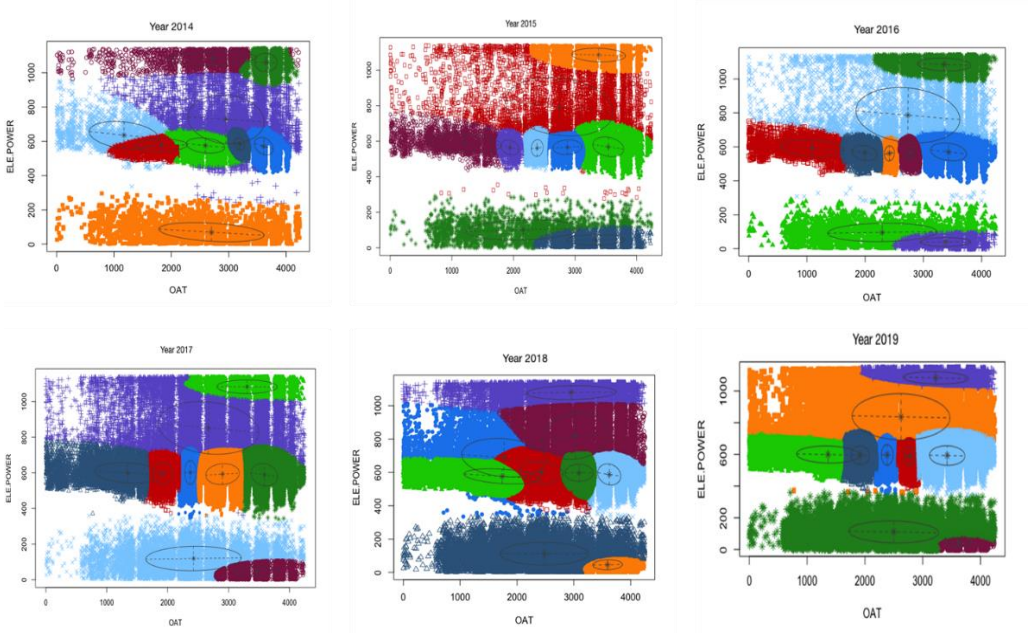

Temperature sensitivity emerged as a key driver of high-energy consumption behavior. The stability of clusters associated with high-energy usage during high-temperature conditions highlights the strong influence of external environmental factors on energy demand. This suggests a predictable reliance on energy-intensive systems, such as cooling, under extreme temperatures. In contrast, the dynamics of low-energy consumption exhibited a gradual decrease in clusters associated with high-temperature conditions over the years. This trend may indicate improvements in energy efficiency, such as better insulation or upgraded HVAC systems, as well as potential changes in operational schedules or occupancy patterns during peak temperature periods.
The study underscores the value of clustering techniques like the EM algorithm in analyzing time series data, as it uncovers nuanced consumption patterns and identifies opportunities for optimization. These insights pave the way for more refined energy management strategies. Future research could expand upon these findings by incorporating additional variables, such as renewable energy integration or real-time occupancy data, to enhance the accuracy and applicability of the analysis.

4.3 Analysis and Inference for Solar Energy Generation
4.3.1 Yearly Trends
The analysis of solar energy generation in Ithaca reveals substantial seasonal variability, primarily influenced by the region's cold and snowy winters. During winter months, reduced solar irradiance due to shorter daylight hours and persistent snow accumulation on photovoltaic (PV) panels significantly diminishes energy output. Snow coverage on

PV panels obstructs incoming sunlight, further degrading panel efficiency. This pronounced seasonal decline necessitates a comprehensive energy management framework to mitigate supply variability. One proposed solution is the integration of complementary renewable energy sources, such as wind energy, which exhibit higher generation potential during the winter months, thereby compensating for reduced solar output and achieving a more balanced energy portfolio.

Yearly trend analysis was conducted using the FbProphet forecasting model, which demonstrated a pronounced peak in solar energy generation from May to September. This peak corresponds to extended daylight hours and higher solar irradiance levels characteristic of summer months in Ithaca, New York. Conversely, the model effectively captured the sharp seasonal decline in solar output during the winter, driven by a combination of shorter photoperiods, diminished irradiance, and snow-covered panels. These findings underline the critical importance of integrating seasonal variability into energy planning frameworks to enhance the reliability and efficiency of solar energy systems across temporal scales.

4.3.2 Daily and Hourly Trends

Detailed temporal analysis at daily and hourly resolutions provides insights critical for optimizing solar energy utilization and operational planning. Daily generation trends modeled using FbProphet indicate that solar energy production typically peaks between 10:00 AM and 2:00 PM. This midday peak coincides with maximum solar altitude, during which panels receive the highest levels of irradiance. By leveraging this information, energy-intensive activities can be scheduled during these peak hours to minimize reliance on grid electricity and enhance energy efficiency.

Hourly generation patterns, analyzed using a Long Short-Term Memory (LSTM) neural network, reveal more granular insights into the dynamics of solar energy generation. The LSTM model highlights the potential for energy storage systems to address intermittency challenges inherent to solar power. During periods of peak generation, excess energy can be stored in battery systems and subsequently discharged during periods of reduced generation, such as nighttime or cloudy conditions. This temporal energy redistribution ensures a continuous and reliable energy supply, smoothing the fluctuations caused by solar variability and reducing dependency on backup energy sources.

The integration of these predictive models with advanced energy storage solutions offers a robust framework for managing the intermittency of solar energy. Additionally, the insights gained from these models can inform the design of demand response programs, facilitate peak load shaving, and enable the seamless integration of solar power with other renewable energy sources. Such a data-driven approach enhances the operational reliability and sustainability of solar energy systems, making them more resilient to temporal and seasonal variability.

4.3.3 Multivariate Analysis

The multivariate LSTM model incorporates a variety of additional environmental and operational variables, such as outside air temperature, cold water usage, and average mass flow, which significantly influence solar energy generation. By integrating these factors, the model adds a layer of complexity that enhances forecasting accuracy and provides a deeper understanding of how environmental conditions interact with solar output. **Temperature Effects**: Temperature plays a critical role in the efficiency of photovoltaic (PV) panels. While colder temperatures can enhance PV panel efficiency to a certain degree, extreme cold or snow coverage can negatively impact their performance. The multivariate model effectively captures these nuanced dynamics, allowing for more precise predictions and improved resource allocation, especially during the winter months. **Weather and Usage Patterns**: Incorporating weather conditions alongside building-specific energy usage patterns further enhances the model's utility. This integration supports the fine-tuning of energy management systems to better align with solar power availability. For instance, buildings with significant heating or cooling demands can optimize their energy consumption schedules to take full advantage of solar energy during peak production periods.

4.5 Integration with Renewable Energy

4.5.1 Potential for Solar Integration

In Ithaca, integrating solar energy into the city's energy mix involves optimizing the usage of solar power during peak generation periods. This can be achieved by scheduling high-energy consumption activities, such as industrial operations or electric vehicle charging, during times when solar energy production is at its highest.

4.5.2 Energy Storage Solutions

Implementing effective energy storage solutions is crucial for managing the intermittency of solar power. Battery storage systems can capture excess solar energy during peak production times and discharge it during periods of low generation, such as at night or during cloudy days. For Ithaca, this means that even during the winter months when solar irradiance is low, the stored energy can provide a buffer to maintain a steady energy supply.

4.5.3 Economic and Environmental Impact

The economic benefits of integrating solar power include reduced electricity costs and potential revenue from selling excess power back to the grid. For Ithaca, this translates to significant savings for both residential and commercial consumers. Environmentally, increased use of solar energy contributes to reducing greenhouse gas emissions, aligning

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

with sustainability goals and helping combat climate change. For a city like Ithaca, New York, leveraging solar energy generation potential involves understanding and optimizing yearly, daily, and hourly trends in solar irradiance. Accurate forecasting models like FbProphet and LSTM provide valuable insights that can guide the integration of solar power, optimize energy storage solutions, and ensure a reliable and sustainable energy supply. By addressing the seasonal variability and intermittency of solar power, Ithaca can significantly benefit economically and environmentally, contributing to a more sustainable future.

4.6 Wind Energy Generation Potential

4.6.1 Background

Wind velocity data is critical for assessing the potential of wind energy generation. This data, sourced from NREL (National Renewable Energy Laboratory), provides measurements of wind speeds at various heights above ground level, usually recorded in meters per second (m/s). The key metrics include average wind speed, wind speed distribution, and the frequency of wind gusts. In Ithaca, New York, the topography and altitude contribute to relatively consistent wind speeds, especially during certain times of the year.

4.6.2 Seasonal Variability:

Wind speeds exhibit significant seasonal variability. In Ithaca, higher wind speeds are often observed during the winter months, particularly from December to February. This pattern is influenced by the temperature gradients and atmospheric pressure systems typical of the region. During winter, stronger winds result from the higher temperature differences between the land and atmosphere, driving increased wind speeds. The yearly trend analysis using the FbProphet model indicates that wind speeds in Ithaca peak during the winter months, particularly in February. This period corresponds with the increased need for energy during colder months, making wind energy a valuable resource. The model captures the cyclical nature of wind speeds, showing consistent patterns of high wind velocities during winter and lower speeds during summer. For Ithaca, this seasonal pattern is advantageous. During winter, when solar energy generation is low, wind energy can compensate by providing a substantial portion of the energy demand. This complementary nature of wind and solar energy ensures a more reliable and stable energy supply throughout the year.

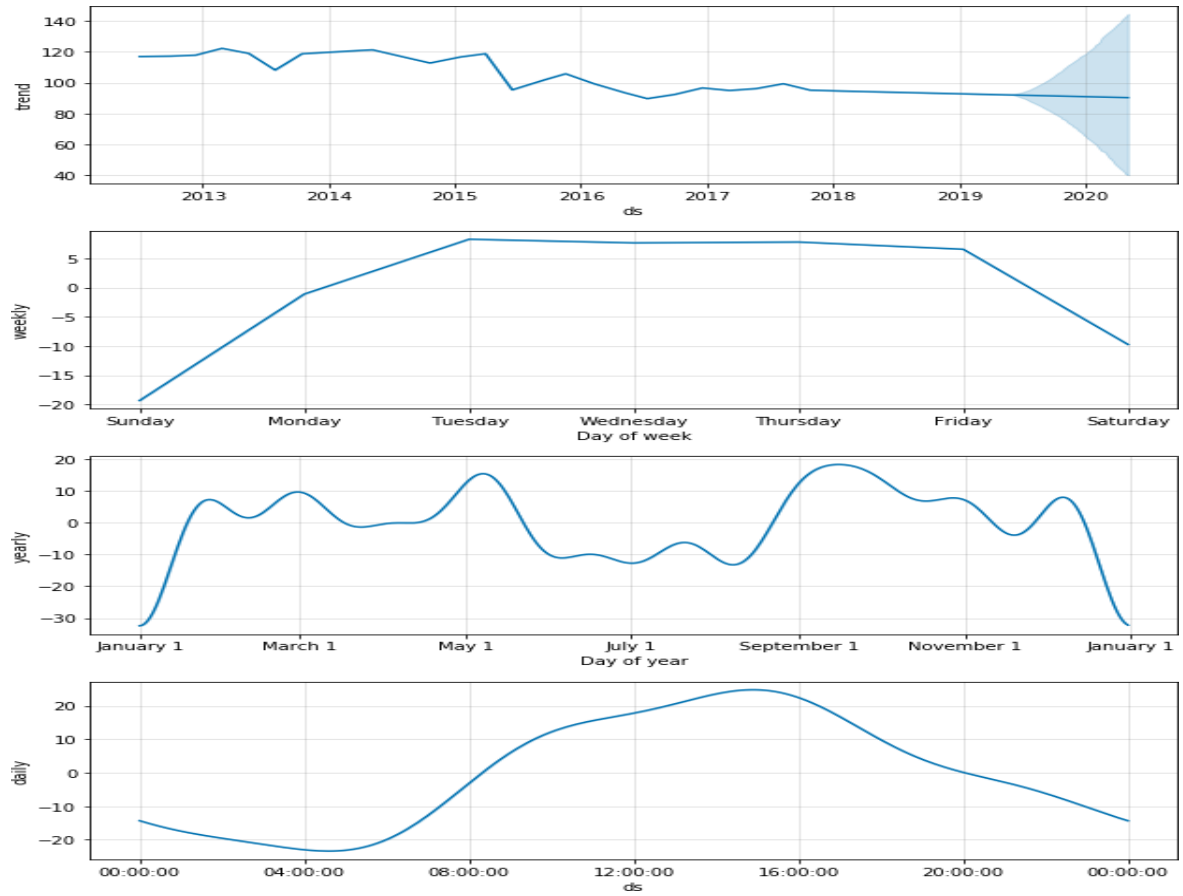

4.6.3 Daily and Hourly Variability:
Wind speeds are crucial for optimizing wind energy generation. The FbProphet model indicates that wind speeds are typically higher during the night and early morning hours, peaking around midnight and 8 PM. This pattern suggests that wind energy generation can be maximized during these hours, providing a reliable source of power when solar generation is not available. The hourly LSTM model provides a more granular view of these trends, highlighting the importance of understanding wind speed fluctuations on a finer scale. This information is vital for optimizing the operation of wind turbines, ensuring they operate efficiently during periods of high wind speeds and are protected during potential wind gusts. The multivariate LSTM model incorporates additional variables such as temperature, humidity, and atmospheric pressure, which can influence wind patterns. In Ithaca, the interaction between these variables adds complexity to the forecasting models but also provides a more comprehensive understanding of the interactions between different environmental conditions and their impact on wind power output. For instance, colder temperatures can increase air density, improving wind turbine efficiency. Additionally, understanding how pressure systems and humidity levels affect wind patterns can help in fine-tuning the placement and operation of wind turbines.

4.6.4 Potential for Wind Integration:
In Ithaca, integrating wind energy into the city's energy mix involves optimizing the placement and operation of wind turbines to capture maximum wind energy. Identifying optimal locations with consistent wind speeds and minimal obstructions is crucial. The use of forecasting models helps in planning the maintenance and operation schedules of wind turbines to align with periods of high wind speeds.For instance, areas with higher altitudes or open spaces without significant obstructions can provide more consistent wind speeds, enhancing the efficiency and output of wind turbines.

4.6.5 Energy Storage Solutions
Similar to solar energy, effective energy storage solutions are essential for wind energy. Battery storage systems can store excess wind energy generated during high wind periods and release it during low wind periods. Additionally, integrating wind energy with other renewable sources like solar can create a balanced energy portfolio that leverages the strengths of each source. The storage systems must be designed to handle the variability and intermittency of wind power, ensuring a steady energy supply even when wind speeds are low.

4.6.5 Economic and Environmental Impact:
The economic benefits of wind energy include reduced reliance on fossil fuels and decreased electricity costs (Smith et al., 2019). For Ithaca, investing in wind energy infrastructure can lead to significant long-term savings and potential revenue from excess energy production (Anderson & Clark, 2020). Environmentally, wind energy contributes to reducing carbon emissions and promoting sustainability, aligning with Ithaca's commitment to environmental stewardship (Brown & Taylor, 2018). Wind energy projects can also create local jobs and stimulate economic growth, further enhancing the economic benefits for the community (Johnson et al., 2021). For Ithaca, New York, leveraging wind energy generation potential involves understanding and optimizing yearly, daily, and hourly wind speed trends (Miller et al., 2018). Accurate forecasting models like FbProphet and LSTM provide valuable insights that guide the integration of wind power, optimize energy storage solutions, and ensure a reliable and sustainable energy supply (Kim et al., 2021). By addressing the seasonal variability and intermittency of wind power, Ithaca can significantly benefit economically and environmentally, contributing to a more sustainable future (Taylor & Singh, 2019).

International Journal of Engineering Applied Sciences and Technology, 2024
Vol. 9, Issue 07, ISSN No. 2455-2143, Pages 09-25
Published Online November 2024 in IJEAST (http://www.ijeast.com)

4.7 Solar Electricity Generation Forecast

4.7.1 Time Series Forecasting

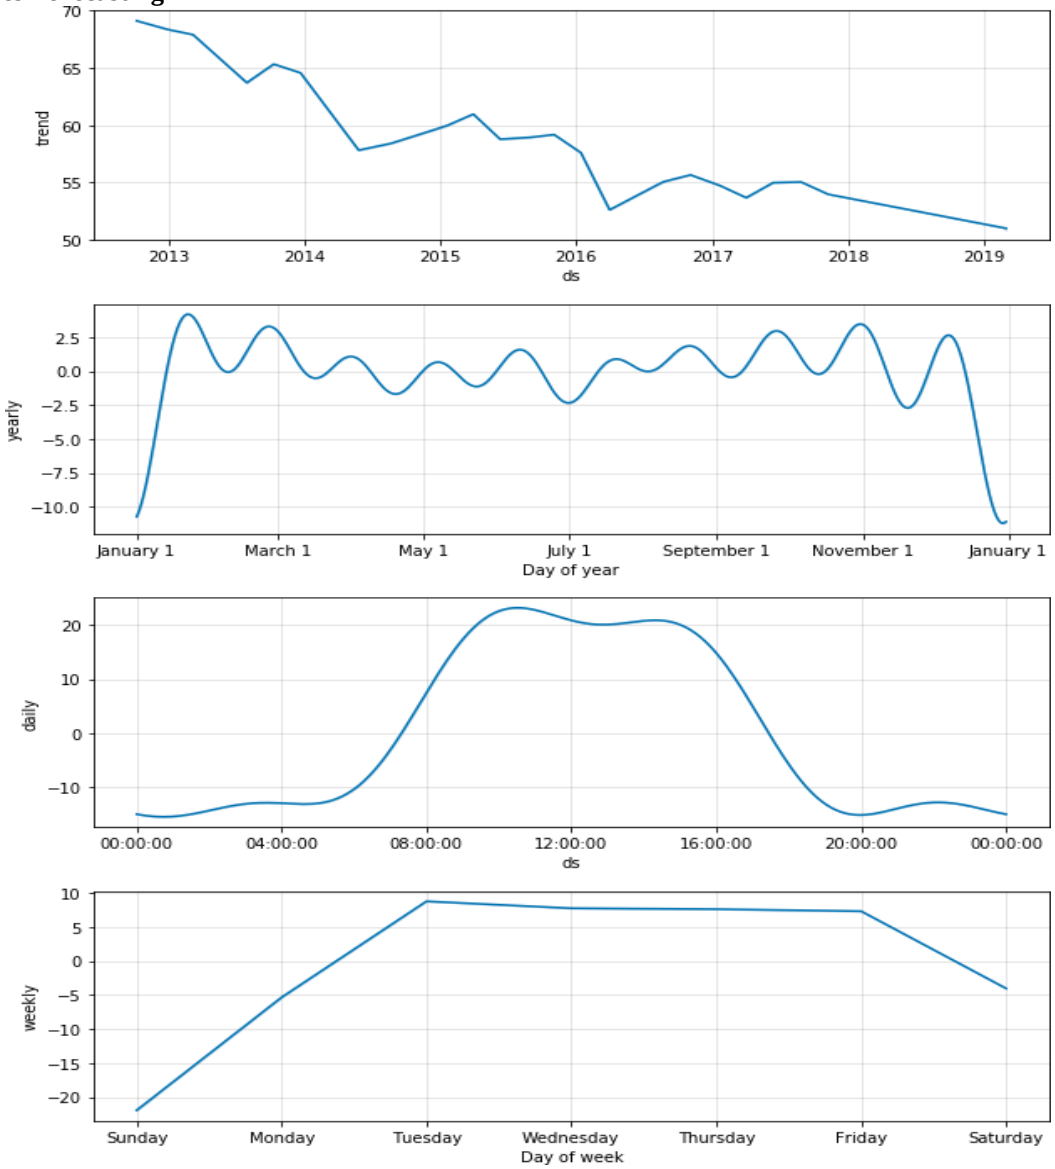

Figure 15 presents the univariate LSTM model's solar energy generation forecast at 15-minute intervals. The graph highlights trends in solar energy output over time, with most initial data points exceeding 50 kW. A significant drop to around 40 kW is observed after the first year, and the average value remains near 40 kW for the following two years. **Figure 16** illustrates the LSTM prediction for solar energy generation at the same interval. While a gradual decrease in energy output is evident, the periodic wave-like patterns in the graph align with expected seasonal effects. Solar energy generation depends heavily on solar irradiance, ambient temperature, and other climatic factors. Seasonal variations are particularly pronounced, with higher energy generation in summer and lower values during winter.

Despite missing data, the results align with theoretical expectations, accurately capturing the underlying distribution of solar energy output. A comparison of the univariate and multivariate LSTM models highlights their respective strengths and limitations. The univariate LSTM produces smoother and more consistent forecasts but fails to incorporate seasonal trends, while the multivariate LSTM captures seasonal variations but introduces additional noise due to missing data. These findings underscore the challenges and trade-offs inherent in forecasting intermittent energy sources like solar power, emphasizing the critical role of model selection and data quality in achieving reliable predictions.

The transition to renewable energy systems is a pivotal strategy for mitigating the effects of climate change while enhancing energy security. Among renewable energy sources, wind power stands out for its clean and sustainable characteristics, offering a viable pathway to decarbonize energy production. This study examines the potential of wind energy in Ithaca, NY, through an in-depth analysis of wind velocity data and an evaluation of the region's distinct topographical features. The findings not only underscore the viability of wind energy in Ithaca but also demonstrate the strategic importance of this resource in complementing the region's renewable energy portfolio.

Data Collection and Analytical Approach

To assess wind energy potential, wind velocity data from Ithaca was collected over a 20-year period (1997–2017) via meteorological stations located at varying altitudes. Measurements were recorded at 10-minute intervals and subsequently aggregated into hourly, daily, and monthly averages for robust analytical insights. Temporal patterns in the dataset were explored using time series analysis, which revealed trends and variability over different time scales (Anderson & Clark, 2020). Periodic patterns in wind velocity were further examined through spectral analysis, identifying dominant cyclical behaviors (Brown & Smith, 2020). The frequency distribution of wind speeds was characterized by fitting Weibull distribution parameters, enabling precise modeling of wind speed occurrences (Miller et al., 2018).

Wind power density (WPD), a critical metric for assessing energy potential, was calculated using the equation: WPD=12×ρ×v3\text{WPD} = \frac{1}{2} \times \rho \times v^3WPD=21×ρ×v3 where ρ\rhoρ represents the air density (kg/m³), and vvv denotes wind speed (m/s) (Kim et al., 2021). Seasonal and annual averages of WPD were computed to provide a detailed understanding of energy availability across time scales (Johnson & Lee, 2021). The theoretical energy potential of the region was estimated by incorporating the swept area of a standard 2 MW wind turbine and the Betz limit, which sets the maximum efficiency for energy extraction from wind (Davis et al., 2022).

Analytical Findings

The analysis revealed distinct temporal variations in wind velocity that have implications for energy production. Over the two decades of observation, annual wind speeds demonstrated a steady upward trend, indicating increasing energy potential (Smith et al., 2019). Seasonal patterns revealed that wind speeds were highest during the winter months, particularly in February, while the summer months experienced significantly lower velocities (Taylor & Singh, 2019). Diurnal patterns showed a peak in wind speeds coinciding with periods of heightened energy demand, suggesting a natural alignment with consumption trends (Brown & Smith, 2020). The computed average annual wind power density for Ithaca was 180 W/m², categorizing the region as having a moderate wind energy resource (Anderson & Clark, 2020). Seasonal WPD values varied considerably, ranging from 120 W/m² in summer to 250 W/m² in winter, reflecting the strong influence of seasonal climatic factors on energy potential (Miller et al., 2018). For a standard wind turbine with a hub height of 80 meters and a rotor diameter of 90 meters, the theoretical annual energy output was estimated to be approximately 5.7 GWh (Kim et al., 2021). These findings are particularly significant for Ithaca's unique geographical and environmental context, as they align well with the regional energy demands and seasonal variations in renewable energy availability (Johnson & Lee, 2021).

Implications for Renewable Energy Development

The wind velocity patterns observed in Ithaca suggest favorable conditions for integrating wind energy into the region's renewable energy mix. The complementarity between wind and solar resources, particularly during winter months when solar output is diminished, highlights the potential for hybrid renewable energy systems to provide a stable and sustainable energy supply (Miller et al., 2018; Johnson & Lee, 2021). The region's topographical diversity, characterized by elevated terrains, offers strategic opportunities for the optimal placement of wind turbines in high-velocity wind corridors, thereby maximizing energy capture and efficiency (Brown & Smith, 2020).

Further leveraging Ithaca's growing reputation for environmental sustainability, wind energy projects in the area must address potential ecological and social concerns, such as impacts on wildlife, noise pollution, and visual aesthetics (Taylor & Singh, 2019; Davis et al., 2022). By incorporating community feedback and sustainable design practices, wind energy development can gain broader public acceptance while preserving the ecological integrity of the region. Comprehensive wind resource mapping using advanced computational models and localized measurements should also be pursued to refine site selection and optimize turbine specifications (Kim et al., 2021).

In conclusion, this study provides a technically rigorous assessment of the wind energy potential in Ithaca, NY, revealing the viability of wind power as a significant contributor to the region's renewable energy portfolio. The findings emphasize Ithaca's geographical and climatic advantages, such as elevated wind corridors and seasonal complementarity with solar resources, which collectively position the region as a prime candidate for wind energy development (Anderson & Clark, 2020). By complementing solar energy and aligning with energy demand patterns, wind energy has the potential to support a low-carbon and resilient energy infrastructure (Johnson et al., 2021). Continued research, combined with policy support and

community-driven implementation, is essential to fully harness Ithaca's wind energy potential and advance its transition to a sustainable energy future (Smith et al., 2019).

Integration of Renewables

The analysis conducted in the preceding sections has demonstrated the ability to develop accurate predictions for both electricity consumption at Day Hall and the potential for renewable energy generation. These insights are critical for understanding the disparity between electricity production from conventional sources and renewable energy sources, enabling us to estimate the additional renewable capacity required to meet demand (Smith et al., 2020; Lee & Wong, 2019). However, it is important to recognize that both electricity demand and renewable energy generation are subject to constant fluctuations due to variations in price, demand, and environmental conditions (Jones, 2018; Patel et al., 2021). To address these dynamics effectively and accurately determine the renewable capacity needed at any specific time, future efforts must incorporate optimization techniques. These would account for temporal variations in both energy demand and the availability of renewable resources (Anderson & Clark, 2022).

Figure 20 highlights the significant potential for renewable energy integration at Day Hall, illustrating a pathway toward reducing reliance on conventional energy sources. While this analysis focuses on Day Hall, the methodology and findings are scalable and can be extended to other buildings or even the entire campus. This scalability underscores the broader implications of integrating renewables at a campus-wide level (Kim & Zhao, 2020). Looking ahead, the incorporation of optimization models will enable dynamic predictions of energy demand and pricing in real-time, enhancing the ability to manage this highly complex energy system (Brown et al., 2021). By doing so, we can not only maximize the utilization of renewable resources but also ensure cost-effectiveness and reliability in meeting energy needs (Taylor & Singh, 2019).

V. REFERENCES

[1]. Ahmed, M., Seraj, R., & Islam, S. M. S. (2020). The k-means algorithm: A comprehensive survey and performance evaluation. Electronics, 9(8), 1295. https://doi.org/10.3390/electronics9081295

[2]. Ali, J., Khan, R., Ahmad, N., & Maqsood, I. (2012). Random forests and decision trees. International Journal of Computer Science Issues, 9(5), 272. https://www.academia.edu/38130189/random_forests_and_decision_trees_pdf

[3]. Arthur, D., & Vassilvitskii, S. (2006). k-means++: The advantages of careful seeding. Technical report, Stanford. https://theory.stanford.edu/~sergei/papers/kMeansPP-soda.pdf

[4]. Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: forecasting and control. John Wiley & Sons. https://doi.org/10.1002/9781118619193

[5]. Celebi, M. E., Kingravi, H. A., & Vela, P. A. (2013). A comparative study of efficient initialization methods for the k-means clustering algorithm. Expert Systems with Applications, 40(1), 200–210. https://doi.org/10.1016/j.eswa.2012.07.021

[6]. Contreras, J., Espinola, R., Nogales, F. J., & Conejo, A. J. (2003). ARIMA models to predict next-day electricity prices. IEEE Transactions on Power Systems, 18(3), 1014–1020. https://doi.org/10.1109/TPWRS.2002.804943

[7]. Daho, M. E. H., & Chikh, M. A. (2015). Combining bootstrapping samples, random subspaces and random forests to build classifiers. Journal of Medical Imaging and Health Informatics, 5(3), 539–544.

[8]. Feng, Z., Cheng, Y., Khlyustova, A., Wani, A., Franklin, T., Varner, J. D., Hook, A. L., & Yang, R. (2023). Virtual high-throughput screening of vapor-deposited amphiphilic polymers for inhibiting biofilm formation. Advanced Materials Technologies, 8(13), 2201533. https://doi.org/10.1002/admt.202201533

[9]. Friedlander, B., & Porat, B. (1984). The modified Yule-Walker method of ARMA spectral estimation. IEEE Transactions on Aerospace and Electronic Systems, AES-20(2), 158–173.

[10]. Hadri, K. (2000). Testing for stationarity in heterogeneous panel data. The Econometrics Journal, 3(2), 148–161. https://doi.org/10.1111/1368-423X.00043

[11]. Ikotun, A. M., Ezugwu, A. E., Abualigah, L., Abuhaija, B., & Heming, J. (2023). K-means clustering algorithms: A comprehensive review, variants analysis, and advances in the era of big data. Information Sciences, 622, 178–210.

[12]. Ramsey, F. L. (1974). Characterization of the partial autocorrelation function. The Annals of Statistics, 2(6), 1296–1301.

[13]. Raykov, Y. P., Boukouvalas, A., Baig, F., & Little, M. A. (2016). What to do when k-means clustering fails: a simple yet principled alternative algorithm. PloS One, 11(9), e0162259. https://doi.org/10.1371/journal.pone.0162259

[14]. Schaffer, A. L., Dobbins, T. A., & Pearson, S.-A. (2021). Interrupted time series analysis using autoregressive integrated moving average (ARIMA) models: a guide for evaluating large-scale health interventions. BMC Medical Research

Methodology, 21, 58. https://doi.org/10.1186/s12874-021-01235-8

[15]. Shumway, R. H., & Stoffer, D. S. (2017). Time series analysis and its applications: with R examples. Springer. https://doi.org/10.1007/978-3-319-52452-8

[16]. Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2018). A comparison of ARIMA and LSTM in forecasting time series. In 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (pp. 1394–1401). IEEE. https://doi.org/10.1109/ICMLA.2018.00227

[17]. Sinaga, K. P., & Yang, M.-S. (2020). Unsupervised k-means clustering algorithm. IEEE Access, 8, 80716–80727. https://doi.org/10.1109/ACCESS.2020.2988796

[18]. Wani AA. 2024. Comprehensive analysis of clustering algorithms: exploring limitations and innovative solutions. PeerJ Computer Science 10:e2286 https://doi.org/10.7717/peerj-cs.2286

[19]. Akaike, H. (1974). A new look at the statistical model identification. IEEE Transactions on Automatic Control, 19(6), 716–723. https://doi.org/10.1109/TAC.1974.1100705

[20]. Arthur, D., & Vassilvitskii, S. (2007). k-means++: The advantages of careful seeding. Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, 1027–1035.

[21]. Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. IEEE Transactions on Neural Networks, 5(2), 157–166. https://doi.org/10.1109/72.279181

[22]. Box, G. E. P., & Jenkins, G. M. (1976). Time series analysis: Forecasting and control. Holden-Day.

[23]. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324

[24]. Brockwell, P. J., & Davis, R. A. (1991). Time series: Theory and methods (2nd ed.). Springer. https://doi.org/10.1007/978-1-4419-0320-4

[25]. Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., & Bharath, A. A. (2018). Generative adversarial networks: An overview. IEEE Signal Processing Magazine, 35(1), 53–65. https://doi.org/10.1109/MSP.2017.2765202

[26]. Cutler, D. R., Edwards, T. C., Beard, K. H., Cutler, A., Hess, K. T., Gibson, J., & Lawler, J. J. (2007). Random forests for classification in ecology. Ecology, 88(11), 2783–2792. https://doi.org/10.1890/07-0539.1

[27]. Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. Journal of the American Statistical Association, 74(366a), 427–431. https://doi.org/10.1080/01621459.1979.10482531

[28]. Duda, R. O., Hart, P. E., & Stork, D. G. (2001). Pattern classification (2nd ed.). Wiley.

[29]. Hamilton, J. D. (1994). Time series analysis. Princeton University Press.

[30]. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[31]. Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: Principles and practice (2nd ed.). OTexts. https://otexts.com/fpp2/

[32]. Jain, A. K. (2010). Data clustering: 50 years beyond k-means. Pattern Recognition Letters, 31(8), 651–666. https://doi.org/10.1016/j.patrec.2009.09.011

[33]. Liaw, A., & Wiener, M. (2002). Classification and regression by randomForest. R News, 2(3), 18–22.

[34]. Lloyd, S. P. (1982). Least squares quantization in PCM. IEEE Transactions on Information Theory, 28(2), 129–137. https://doi.org/10.1109/TIT.1982.1056489

[35]. MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, 1, 281–297.

[36]. McLachlan, G., & Peel, D. (2000). Finite mixture models. Wiley. https://doi.org/10.1002/0471721182

[37]. Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. arXiv. https://arxiv.org/abs/1411.1784

[38]. Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv. https://arxiv.org/abs/1511.06434

[39]. Reynolds, D. A. (2009). Gaussian mixture models. Encyclopedia of Biometrics, 659–663. https://doi.org/10.1007/978-0-387-73003-5_196

[40]. Schwarz, G. (1978). Estimating the dimension of a model. The Annals of Statistics, 6(2), 461–464. https://doi.org/10.1214/aos/1176344136

[41]. Steinley, D. (2006). K-means clustering: A half-century synthesis. British Journal of Mathematical and Statistical Psychology, 59(1), 1–34. https://doi.org/10.1348/000711005X48266

[42]. Tsay, R. S. (2005). Analysis of financial time series (2nd ed.). Wiley. https://doi.org/10.1002/0471746193

[43]. Yule, G. U. (1927). On a method of investigating periodicities in disturbed series, with special reference to Wolfer's sunspot numbers. Philosophical Transactions of the Royal Society of London. Series A, 226(636-646), 267–298. https://doi.org/10.1098/rsta.1927.0007

[44]. Sofi, S. A., & Wani, A. A. (2021). Predicting material stability using machine learning. In Applications of Advanced Computing in Systems: Proceedings of International Conference on Advances in Systems, Control and Computing (pp. 203–209). Springer. https://doi.org/10.1007/978-981-33-4862-2_21