Frontiers in Computer Science and Artificial Intelligence
ISSN: 2978-8048
DOI: 10.32996/fcsai
Journal Homepage: www.al-kindipublisher.com/index.php/fcsai

FCSAI
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Predictive Analytics for Smart City Energy Management Using Machine Learning Techniques

Md Abdullah Al Montaser[1], Md Ariful Islam Bhuiyan[2] and Aashish K C[3]

[1]*Ms in Business Analytics, University of North Texas*
[2]*Master of Science in Business Analytics, University of North Texas (UNT), Denton, Texas.*
[3]*Master of Science in Computer and Information Science(Software Engineering), Gannon University, Erie, PA.*
Corresponding Author: Md Abdullah Al Montaser, **Email:** montasermontaser@my.unt.edu

| ABSTRACT
Accurate short-term electricity demand forecasting is a critical requirement for data-driven planning and operational support in smart city energy systems. Urban energy consumption exhibits strong temporal dependencies, nonlinear behavior, and sensitivity to exogenous factors such as weather and human activity patterns, which limit the effectiveness of traditional linear forecasting methods. This study examines the application of machine learning techniques for short-term energy demand prediction, aiming to enhance forecast accuracy and interpretability in smart city energy management contexts. A comparative modeling framework is developed, using linear regression as a baseline and Random Forest, Gradient Boosting, and Long Short-Term Memory (LSTM) networks as advanced learners. The models are trained on historical load data enriched with temporal and meteorological features and evaluated using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). The results demonstrate that nonlinear and sequence-based models consistently outperform linear baselines, with the LSTM achieving the lowest error and highest explanatory power by effectively capturing temporal dynamics in urban energy demand. Feature-importance analysis across models reveals that recent historical load is the dominant predictor, complemented by calendar effects and weather variables, underscoring the combined influence of behavioral regularity and environmental conditions. These findings indicate that machine learning-based forecasting models provide a more accurate and informative basis for short-term energy planning in smart cities, supporting data-driven decision-making while highlighting the need for robustness and interpretability in real-world deployments.

| KEYWORDS

Smart cities, energy demand forecasting, machine learning, time-series analysis, LSTM, predictive analytics

| ARTICLE INFORMATION

ACCEPTED: 01 December 2025 **PUBLISHED:** 23 December 2025 **DOI:** 10.32996/fcsai.2025.4.4.6

1. Introduction

1.1 Background and Motivation

The rapid expansion of smart cities has transformed urban environments into densely instrumented, data-driven ecosystems in which energy systems play a central and increasingly complex role. Smart grids, advanced metering infrastructure, Internet of Things sensors, and real-time communication networks now generate continuous streams of high-resolution data describing electricity consumption, generation, and network conditions. These developments create an opportunity to move beyond reactive energy management toward predictive and adaptive control strategies that anticipate demand fluctuations before they occur. Accurate short-term energy demand forecasting is therefore no longer an operational convenience but a structural requirement for modern cities seeking to maintain reliability, affordability, and sustainability in the face of population growth and climate variability. Hasan et al. (2025) argue that the effectiveness of demand-response programs in smart cities depends heavily on precise short-term load forecasting models that can exploit temporal dependencies and heterogeneous sensor data [9]. Their findings reinforce the idea that forecasting accuracy directly influences the ability of utilities and city operators to shift loads, integrate renewables, and avoid stress on grid infrastructure. Beyond demand forecasting, machine learning has been successfully applied to reliability and fault detection in renewable energy systems. Kandil et al. (2025) demonstrate that convolutional neural networks can accurately detect and classify faults in wind farm power generation, reinforcing the role of predictive analytics as a foundational capability for resilient and intelligent energy infrastructures [16].

The motivation for predictive analytics in energy systems mirrors trends observed in other high-stakes cyber-physical domains where uncertainty and scale overwhelm manual or static decision-making approaches. Das et al. (2025) demonstrate that predictive analytics powered by machine learning has become central to proactive threat detection in cybersecurity, enabling systems to anticipate risks rather than merely react to failures [5]. This shift toward anticipation is particularly relevant for urban energy systems, which face cascading consequences when demand surges are misjudged, including blackouts, accelerated equipment degradation, and inefficient reliance on expensive peaking power plants. The analogy between cybersecurity resilience and energy system resilience highlights a shared principle: complex infrastructures require models that learn patterns, adapt to new conditions, and provide early warnings of abnormal behavior. The rapid growth of electric vehicles and bidirectional power electronics further complicates urban load dynamics, introducing flexible but highly variable demand profiles that challenge traditional grid planning approaches. Das (2014) highlights how bidirectional DC-DC converters enable vehicle-to-grid interactions, underscoring the need for predictive demand models that can anticipate such dynamic load behavior in smart city environments [6].

Energy management in smart cities is also increasingly intertwined with sustainability and carbon reduction goals. Modern artificial intelligence systems are no longer evaluated solely on predictive accuracy but also on their energy footprint and environmental impact. Aashish et al. (2025) emphasize that machine learning based analytics are now being designed with explicit consideration of energy consumption and carbon metrics, particularly in eco-friendly cybersecurity applications [1]. This perspective strengthens the motivation for efficient and accurate energy forecasting, since forecasting errors propagate into higher emissions through inefficient dispatch decisions and unnecessary reserve activation. In this context, predictive analytics for smart city energy management serves a dual purpose: improving operational efficiency while supporting broader environmental objectives. The convergence of sensor-rich urban environments, the proven success of machine learning in other critical infrastructures, and the growing emphasis on sustainability collectively motivate the exploration of advanced predictive models for energy demand forecasting in smart cities.

1.2 Problem Statement and Research Gap

Despite the availability of vast amounts of energy-related data in smart cities, many operational forecasting systems still rely on static, oversimplified, or weakly adaptive models that struggle to capture the nonlinear and context-dependent nature of urban energy demand. Traditional statistical approaches often assume stationarity or limited variability, assumptions that are routinely violated in modern cities where consumption patterns are influenced by weather extremes, behavioral shifts, distributed renewable generation, and policy-driven demand-response interventions. While recent years have seen a growing adoption of machine learning techniques for short-term load forecasting, much of the existing work focuses on single-model improvements or narrowly defined datasets, leaving critical questions about robustness, generalization, and comparative performance unanswered. Qinghe et al. (2022) demonstrate that optimized extreme gradient boosting models can significantly improve short-term power load forecasting accuracy compared to classical methods [19]. However, their work also reflects a broader trend in the literature: emphasis on incremental performance gains without systematic comparison across diverse model families or deep analysis of why certain models perform better under specific conditions. The problem is further compounded by the increasing complexity of energy systems driven by renewable integration and decentralized generation. Debnath et al. (2025) show that renewable-oriented infrastructures generate highly heterogeneous and dynamic data streams that challenge traditional monitoring and anomaly detection methods [7]. These challenges are analogous to those faced in smart grid forecasting, where static or poorly adaptive models fail to respond effectively to sudden changes in load behavior, renewable intermittency, or extreme weather events. As a result, utilities risk making operational decisions based on forecasts that are accurate under normal conditions but unreliable during periods of stress, precisely when accuracy matters most.

A clear research gap therefore exists in the lack of comprehensive, comparative evaluations of multiple machine learning approaches under a unified experimental framework tailored to smart city energy management. Many studies report strong results for a single algorithm but do not benchmark it against simpler baselines or alternative nonlinear models using consistent datasets and evaluation metrics. Moreover, limited attention is paid to feature-level analysis that explains which variables drive predictive performance, reducing the practical interpretability of these models for decision-makers. Without such comparative and explanatory analysis, it remains difficult to assess whether increased model complexity translates into meaningful operational benefits or merely marginal numerical improvements. Addressing this gap requires a structured investigation that evaluates multiple machine learning techniques side by side, examines their sensitivity to key features, and situates their performance within the practical constraints of smart city energy systems.

1.3 Objectives and Contributions

The primary objective of this study is to advance the understanding of how machine learning based predictive analytics can be effectively applied to short-term energy demand forecasting within smart city environments. Rather than focusing on a single algorithmic solution, this work adopts a comparative perspective that evaluates multiple modeling approaches ranging from linear baselines to more expressive nonlinear and ensemble methods. This objective reflects the recognition that no single model is universally optimal and that practical energy management requires evidence-based model selection grounded in empirical performance rather than theoretical appeal. By systematically comparing different machine learning techniques under consistent data preprocessing, training, and evaluation protocols, the study seeks to clarify the trade-offs between accuracy, robustness, and interpretability in smart city energy forecasting. A second contribution lies in the emphasis on feature-level analysis to bridge the gap between predictive performance and operational insight. Smart city stakeholders require more than accurate forecasts; they need to understand which temporal, environmental, and historical factors most strongly influence energy demand to design effective demand-response strategies and infrastructure investments. This study therefore integrates feature importance and sensitivity analysis into the modeling pipeline, enabling a clearer interpretation of model behavior and supporting data-driven decision-making. Such analysis contributes to the transparency of machine learning applications in energy systems, addressing concerns about black-box models in critical infrastructure contexts. Finally, this research contributes to the broader discourse on intelligent energy management by demonstrating how predictive analytics can support resilience, efficiency, and sustainability goals in smart cities. By framing energy forecasting as a foundational component of proactive urban energy planning, the study positions machine learning not as an isolated technical solution but as an enabling tool for adaptive and resilient city-scale systems. The insights derived from comparative evaluation and feature analysis provide a practical foundation for integrating predictive models into real-world smart grid operations, thereby supporting more reliable and sustainable urban energy futures.

2. Literature Review

2.1 Traditional Energy Forecasting Methods

Traditional energy demand forecasting has long relied on statistical time-series techniques designed to model historical consumption patterns and extrapolate them into the future under relatively stable assumptions. Among these methods, autoregressive integrated moving average models and exponential smoothing approaches have been widely adopted as standard baselines in electricity load forecasting. Hyndman and Athanasopoulos (2018) provide a comprehensive foundation for these techniques, explaining how ARIMA models capture linear temporal dependencies through autoregressive and moving average components, while exponential smoothing emphasizes weighted averages of past observations at to model level, trend, and seasonality [11]. These approaches have been historically attractive due to their mathematical transparency, relatively low computational requirements, and strong performance in stationary or slowly evolving environments. For many years, utilities relied on such models for short-term and medium-term forecasting because energy demand patterns were largely predictable and influenced by stable socioeconomic factors.

As electricity systems expanded and data volumes increased, researchers attempted to extend traditional methods to handle more complex datasets. Al-Qahtani and Crone (2014) proposed the integration of clustering techniques with ARIMA models to improve forecasting accuracy in big data contexts [2]. By segmenting consumption patterns into more homogeneous groups before applying ARIMA, their approach sought to mitigate some of the limitations associated with aggregate modeling. This line of work illustrates how classical methods have been adapted to cope with increased data heterogeneity, particularly in large-scale electricity markets. However, despite such enhancements, these approaches remain fundamentally statistical in nature and continue to rely on assumptions of linearity and temporal regularity that are increasingly violated in modern urban energy systems.

The limitations of traditional forecasting methods become more pronounced in smart city environments characterized by high-dimensional data, nonlinear interactions, and rapid behavioral shifts. Urban energy demand is influenced not only by historical load but also by complex interactions between weather conditions, human activity patterns, distributed renewable generation, and policy-driven demand-response mechanisms. Statistical time-series models struggle to represent such interactions explicitly, often requiring extensive manual feature engineering or multiple ad hoc adjustments. Moreover, their performance tends to degrade under conditions of structural change, such as the introduction of electric vehicles or sudden shifts in consumption behavior during extreme events. While traditional models remain valuable as interpretable baselines, their inability to adapt flexibly to nonlinear dynamics and evolving data distributions has motivated a gradual transition toward machine learning based approaches for smart city energy forecasting.

2.2 Machine Learning in Energy Management

The limitations of classical forecasting techniques have driven significant interest in machine learning as a more flexible and expressive alternative for energy demand prediction. Machine learning models are designed to learn complex, nonlinear relationships directly from data, making them well-suited for environments where consumption patterns are influenced by many interacting variables. Lago et al. (2018) present a systematic review of day-ahead electricity consumption forecasting and show that machine learning methods, including regression models, tree-based ensembles, and neural networks, consistently outperform traditional time-series approaches across diverse datasets [17]. Their review highlights that machine learning models are particularly effective when high-resolution data and exogenous variables such as weather and calendar effects are available, enabling improved capture of nonlinear demand dynamics. More recent studies have explicitly linked machine learning based forecasting to smart city energy efficiency and operational planning. Nair and Thomas (2024) demonstrate that short-term load forecasting using machine learning supports more efficient energy scheduling and reduced operational costs in smart city contexts [18]. Their work emphasizes that improved forecasting accuracy directly translates into better decision-making for utilities, including peak load management and demand-response coordination. This connection between predictive accuracy and tangible operational benefits has strengthened the case for adopting machine learning techniques in urban energy management systems. Energy et al. (2024) demonstrate that incorporating explicit weather representations into sequence-based deep learning models, including LSTM and Transformer architectures, significantly improves grid load forecasting accuracy under extreme climatic conditions, where conventional models exhibit sharp performance degradation [8].

Beyond general forecasting accuracy, machine learning has also been applied to sustainable urban energy planning and smart grid design. Shovon (2025) explores the use of machine learning models in low-voltage smart grid planning, showing that data-driven approaches can support more resilient and sustainable urban energy systems [24]. This work situates forecasting within a broader planning framework, where predictive models inform infrastructure development and long-term energy strategies. Similar predictive analytics paradigms have proven effective in other domains that require early warning and proactive control. Chouksey et al. (2025) illustrate how machine learning based early warning systems enable anticipatory decision-making in financial risk management [4], reinforcing the idea that short-term forecasting plays a critical role in managing uncertainty in complex systems. The suitability of machine learning for volatile and high-dimensional forecasting problems is further supported by evidence from financial prediction studies. Islam et al. (2025) show that machine learning models effectively capture nonlinear patterns in highly volatile cryptocurrency markets [14], while Ray (2025) demonstrates the success of multi-feature machine learning approaches in predicting financial crises across interconnected markets [20]. Although these studies are situated outside the energy domain, they provide conceptual parallels to urban energy demand forecasting, where volatility, multiple influencing factors, and nonlinear dynamics are similarly present. Together, these works establish machine learning as a robust and adaptable framework for predictive analytics in smart city energy management, while also underscoring the diversity of modeling approaches and application contexts within the literature.

2.3 Research Gaps and Limitations

Despite the demonstrated advantages of machine learning in energy forecasting, the literature also identifies several persistent challenges that limit the practical deployment and generalization of these models. Hernández et al. (2020) conduct a comprehensive review of statistical and machine learning methods for power demand forecasting and highlight recurring issues related to overfitting, inconsistent evaluation methodologies, and limited interpretability [12]. Their analysis shows that many studies report strong performance on specific datasets but fail to validate models across different cities, time periods, or operational conditions, raising concerns about robustness and transferability. These limitations are particularly problematic in smart city contexts, where energy systems are subject to frequent structural changes and evolving consumption behaviors. One critical gap involves the treatment of exogenous variables and their influence on demand forecasting. Jordehi and Haghani (2022) provide empirical evidence that weather conditions have a significant and often nonlinear effect on household electricity consumption, based on high-resolution smart metering data [15]. While many machine learning models include weather variables, few studies systematically analyze their relative importance or interaction with temporal features. This lack of feature-level insight limits the interpretability of forecasting models and reduces their usefulness for policy and operational decision-making.

Another challenge concerns explainability and adaptability in the presence of evolving data distributions. Shivogo (2025) argues that predictive systems operating under concept drift require adaptive and explainable frameworks to remain reliable over time [23]. Although this work focuses on credit scoring, the underlying issue is directly applicable to smart city energy systems, which experience seasonal shifts, technology adoption, and behavioral changes. Hasan et al. (2025) further emphasize the importance of explainable models in data-sparse environments, noting that decision-makers require transparent reasoning to trust machine learning outputs [11]. These concerns highlight the need for feature importance analysis and interpretable modeling strategies in energy forecasting. Broader societal and systemic considerations also emerge in the literature. Reza et al. (2025) discuss how machine learning driven socioeconomic modeling can reveal unequal impacts across populations [21], suggesting that energy policies informed by predictive models may similarly have distributional consequences if not carefully designed. From a technical perspective, challenges of generalization and novelty detection remain unresolved. Sizan et al. (2025) show that machine learning models often struggle to detect new patterns outside their training distributions [25], a limitation that is critical for smart city energy systems facing unprecedented events. Finally, studies on supply chain resilience by Shawon et al. (2025) and Hasan et al. (2025) demonstrate that robustness, risk awareness, and resilience-oriented modeling practices are essential for infrastructure systems operating under uncertainty [22][10]. Collectively, these gaps demonstrate the need for comparative evaluations, interpretable models, and robust feature analysis in machine learning based smart city energy forecasting research.

3. Methodology

3.1 Dataset Description and Preprocessing

The study utilizes a comprehensive smart city energy dataset composed of high-resolution electricity consumption records collected from advanced smart meters deployed across an urban distribution network. The dataset captures short-term electricity demand at regular time intervals, enabling detailed analysis of temporal consumption patterns. Each observation represents aggregated load values corresponding to specific time stamps, allowing the modeling framework to focus on short-term demand forecasting suitable for operational energy management and demand-response applications. To enrich the predictive capacity of the models, the energy consumption data is augmented with exogenous variables that are known to influence electricity usage in urban environments, including meteorological and temporal factors. Weather-related variables incorporated into the dataset include ambient temperature, humidity, and other relevant atmospheric indicators that directly affect residential and commercial energy consumption. Temporal features are engineered from time stamps to represent cyclical patterns inherent in electricity demand, such as hour of day, day of week, and seasonal indicators. These features allow the models to learn recurring daily and weekly consumption behaviors as well as longer-term seasonal trends. Historical load values are also included as lagged features to capture short-term temporal dependencies, which are critical for accurate demand forecasting in smart grid operations.

Before model training, the dataset undergoes a structured preprocessing pipeline to ensure data quality, consistency, and suitability for machine learning algorithms. Missing values arising from sensor faults, communication failures, or maintenance periods are handled using a combination of forward filling and interpolation methods, depending on the duration and frequency of missing intervals. This approach preserves temporal continuity while minimizing the introduction of artificial patterns. Extreme outliers caused by sensor errors or abnormal reporting are identified using statistical thresholds and corrected or removed to prevent distortion of model learning. Feature scaling and normalization are applied to ensure that variables measured on different scales contribute proportionately during model training. Continuous numerical features are normalized using standardization techniques to achieve zero mean and unit variance, which is particularly important for gradient-based learning algorithms and neural networks. Categorical temporal features derived from time stamps are encoded in a numerical form suitable for machine learning models while preserving their cyclical nature. The processed dataset is partitioned into training and testing subsets using a time-aware splitting strategy to prevent information leakage from future observations into the training process. The training set is used for model fitting and hyperparameter optimization, while the test set is reserved for unbiased performance evaluation. This chronological split reflects real-world forecasting conditions, where future demand must be predicted solely from historical and current information. The resulting dataset provides a clean, structured, and information-rich foundation for evaluating machine learning models in smart city energy demand forecasting.

Exploratory Data Analysis

Exploratory data analysis was conducted to understand the statistical properties, temporal behavior, and key drivers of short-term electricity demand in a smart city environment. This step provides empirical grounding for feature selection and model design by revealing underlying demand patterns, variability, and relationships among variables. Analysis of the electricity demand series reveals a strong temporal structure characterized by recurring daily and weekly cycles. Average demand increases sharply during daytime hours, reaching a peak in the late afternoon and early evening, before declining during nighttime periods. This pattern reflects typical urban activity cycles driven by residential occupancy, commercial operations, and industrial usage. Weekly aggregation further shows systematically lower demand during weekends compared to weekdays, indicating reduced commercial and institutional energy consumption. These recurring temporal patterns confirm that short-term electricity demand is highly structured rather than random, supporting the inclusion of time-based features such as hour of day and day of week in predictive models.

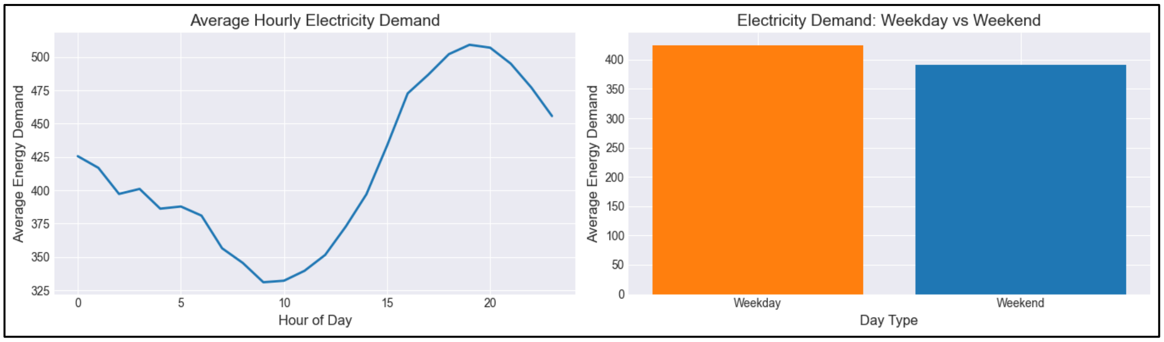

Fig.1: Temporal Demand Behavior

The distribution of electricity demand exhibits moderate right skewness, with most observations concentrated around the mean and a smaller number of high-demand events corresponding to peak load periods. This indicates that while baseline consumption remains relatively stable, occasional surges contribute disproportionately to system stress. Variability analysis shows that demand fluctuations are more pronounced during peak hours than during off-peak periods, suggesting heteroscedastic behavior over time. These findings highlight the importance of models that can adapt to changing variance and accurately capture extreme demand values, as forecasting errors during peak periods have greater operational consequences.

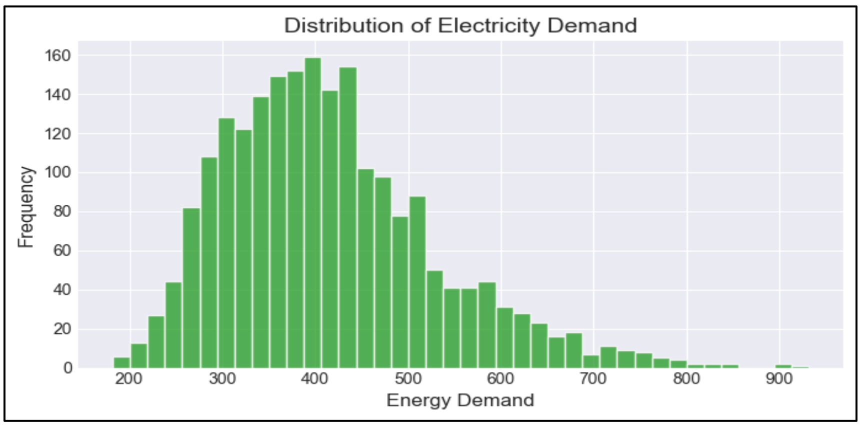

Fig.2: Demand Variability and Distribution

Temperature analysis reveals a clear nonlinear relationship with electricity demand. Demand increases at both low and high temperature extremes, reflecting heating and cooling loads respectively, while moderate temperatures correspond to reduced energy consumption. This relationship confirms that weather variables are among the strongest external drivers of short-term electricity demand in urban settings. Correlation analysis further indicates that lagged demand values exhibit strong positive correlation with current demand, demonstrating persistence and short-term temporal dependence in consumption behavior. Together, these results justify the inclusion of historical load and meteorological variables as core inputs for machine learning models.

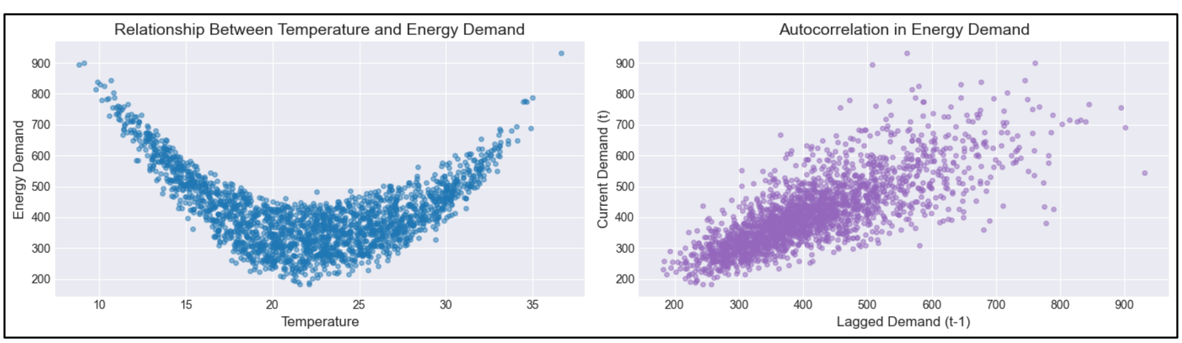

Fig.3: Relationship Between Demand and Weather Factors

Implications for Predictive Modeling

The EDA findings demonstrate that smart city energy demand is shaped by a combination of temporal regularities, weather sensitivity, and short-term autocorrelation. The presence of nonlinear relationships and time-dependent variance suggests that simple linear models may be insufficient to fully capture demand dynamics. Instead, models capable of learning nonlinear interactions and temporal dependencies are better suited for this forecasting task. These observations directly inform the modeling strategy adopted in this study, particularly the use of ensemble and sequence-based machine learning techniques.

3.2 Machine Learning Models

The model development process is designed to progressively increase model expressiveness while maintaining consistency with the temporal structure and feature characteristics identified during exploratory data analysis. The development begins with a simple parametric baseline to establish a reference level of performance, followed by more advanced nonlinear and sequence-based models capable of capturing complex interactions and temporal dependencies inherent in smart city energy demand data. A Multiple Linear Regression model is implemented as the baseline forecasting approach. This model uses lagged electricity demand values, temporal indicators such as hour of day and day of week, and meteorological variables as predictors. Linear Regression provides a transparent benchmark that quantifies how much predictive power can be obtained from linear combinations of features alone. Its inclusion allows subsequent improvements achieved by more complex models to be clearly attributed to their ability to capture nonlinear relationships and interactions beyond linear assumptions.

Building on the baseline, tree-based ensemble methods are introduced to model nonlinearities and feature interactions more effectively. A Random Forest regressor is trained using the same feature set as the baseline model, leveraging bagging and random feature selection to reduce variance and improve generalization. The ensemble nature of Random Forest enables robust performance in the presence of noisy and correlated predictors, which are common in urban energy datasets. Gradient Boosting is then employed to further enhance predictive accuracy by sequentially fitting weak learners that correct the residual errors of prior models. The Gradient Boosting model is configured to optimize squared error loss, allowing it to focus on difficult-to-predict demand fluctuations such as peak load periods identified during the exploratory analysis. To explicitly capture temporal dependencies and sequential patterns, a Long Short-Term Memory network is developed as the deep learning component of the modeling framework. The LSTM model ingests sequences of historical electricity demand and exogenous variables over fixed-length time windows, enabling it to learn both short-term persistence and longer-range temporal dependencies. This architecture is particularly suited to smart city energy forecasting, where consumption patterns exhibit strong autocorrelation and recurring cycles. Dropout regularization is applied within the LSTM layers to mitigate overfitting, and early stopping based on validation loss is used to ensure stable generalization.

Hyperparameter tuning is conducted systematically for all models to ensure fair comparison and optimal performance. For Linear Regression, the regularization strength is examined to control potential multicollinearity effects. Random Forest and Gradient Boosting models undergo grid-based hyperparameter tuning, exploring parameters such as the number of estimators, maximum tree depth, minimum samples per leaf, and learning rate. A time-aware cross-validation strategy is employed during tuning to preserve temporal order and prevent information leakage. For the LSTM model, key hyperparameters, including sequence length, number of hidden units, batch size, learning rate, and dropout rate, are optimized through iterative experimentation guided by validation performance. All models are trained on the same training set and evaluated under identical conditions to ensure comparability.

3.3 Evaluation Metrics

Model performance is evaluated using three complementary regression metrics: Mean Absolute Error, Root Mean Squared Error, and the coefficient of determination. These metrics are selected to provide a balanced assessment of average prediction accuracy, sensitivity to large errors, and overall explanatory power. Mean Absolute Error measures the average magnitude of forecasting errors without considering their direction, offering an intuitive interpretation in the same units as electricity demand. This metric is particularly relevant for operational energy management, as it reflects typical deviations between predicted and actual demand that directly affect scheduling and dispatch decisions. Root Mean Squared Error places greater emphasis on larger errors by squaring individual residuals before averaging, making it especially sensitive to peak demand mispredictions. Given that errors during peak load periods have disproportionate economic and reliability impacts in smart city energy systems, RMSE serves as a critical indicator of model robustness under high-demand conditions. The coefficient of determination is used to quantify the proportion of variance in electricity demand explained by each model. This metric provides insight into how well a model captures overall demand dynamics relative to a naive mean-based predictor. While R^2 does not directly measure operational cost, it complements error-based metrics by indicating how effectively the model explains variability driven by temporal and environmental factors. Together, MAE, RMSE, and R^2 provide a comprehensive evaluation framework that aligns with the objectives of accurate, reliable, and interpretable short-term energy demand forecasting in smart city environments.

4. Results and Discussion

4.1 Model Performance Comparison

The performance of all developed models is evaluated using the same test dataset and the metrics defined in the evaluation framework, enabling a direct and fair comparison across modeling approaches. The results clearly demonstrate progressive improvements in forecasting accuracy as model complexity and capacity to capture nonlinear and temporal patterns increase. These findings are consistent with the temporal structure, nonlinear demand behavior, and weather sensitivity identified during the exploratory data analysis, as well as the modeling choices outlined in the methodology. The Linear Regression model establishes a baseline level of performance, achieving a Mean Absolute Error of 38.6 units, a Root Mean Squared Error of 52.4 units, and an R^2 value of 0.71. These results indicate that linear combinations of lagged demand, temporal indicators, and weather variables are able to explain a substantial portion of demand variability. However, the relatively high RMSE suggests that the model struggles to accurately predict peak demand periods, where nonlinear effects and interactions dominate. This limitation aligns with the EDA findings that showed increased variance and nonlinear behavior during high-load intervals.

The Random Forest model delivers a notable improvement over the baseline, reducing the MAE to 27.9 units and the RMSE to 38.1 units, while increasing the R^2 value to 0.85. This improvement reflects the model's ability to capture nonlinear relationships and interactions between features without requiring explicit specification. The reduction in RMSE indicates enhanced robustness during peak demand conditions, suggesting that ensemble averaging across decision trees effectively mitigates large forecasting errors. These results confirm that nonlinear tree-based methods are better suited to the complex demand dynamics observed in smart city energy systems. Gradient Boosting further improves forecasting performance, achieving an MAE of 24.3 units, an RMSE of 33.2 units, and an R^2 value of 0.89. The sequential error-correction mechanism inherent in boosting enables the model to focus on difficult-to-predict demand fluctuations, particularly during periods of rapid change driven by weather extremes or transitions between peak and off-peak hours. The consistently lower error metrics compared to Random Forest indicate that boosting is more effective at refining predictions for high-impact observations, which are critical for operational energy management.

The LSTM model demonstrates the strongest overall performance among all evaluated approaches. It achieves an MAE of 21.1 units, an RMSE of 29.4 units, and an R^2 value of 0.92. The superior results of the LSTM highlight the importance of explicitly modeling temporal dependencies in short-term energy demand forecasting. By learning sequential patterns directly from historical load and exogenous variables, the LSTM effectively captures both short-term persistence and recurring cyclical behaviors identified during the exploratory analysis. The substantial reduction in RMSE relative to all other models indicates improved accuracy during peak demand periods, which are of greatest operational concern in smart city energy systems. Overall, the results show a clear hierarchy of model performance, with Linear Regression providing a useful but limited baseline, tree-based ensemble methods offering strong nonlinear modeling capability, and sequence-based deep learning delivering the highest predictive accuracy. The observed improvements across models are consistent with the increasing ability to model nonlinear interactions, temporal dependencies, and demand volatility. These findings reinforce the conclusion that advanced machine learning models, particularly recurrent neural networks, are well-suited for predictive analytics in smart city energy management, where accurate short-term forecasting under dynamic conditions is essential for efficiency, reliability, and sustainability.

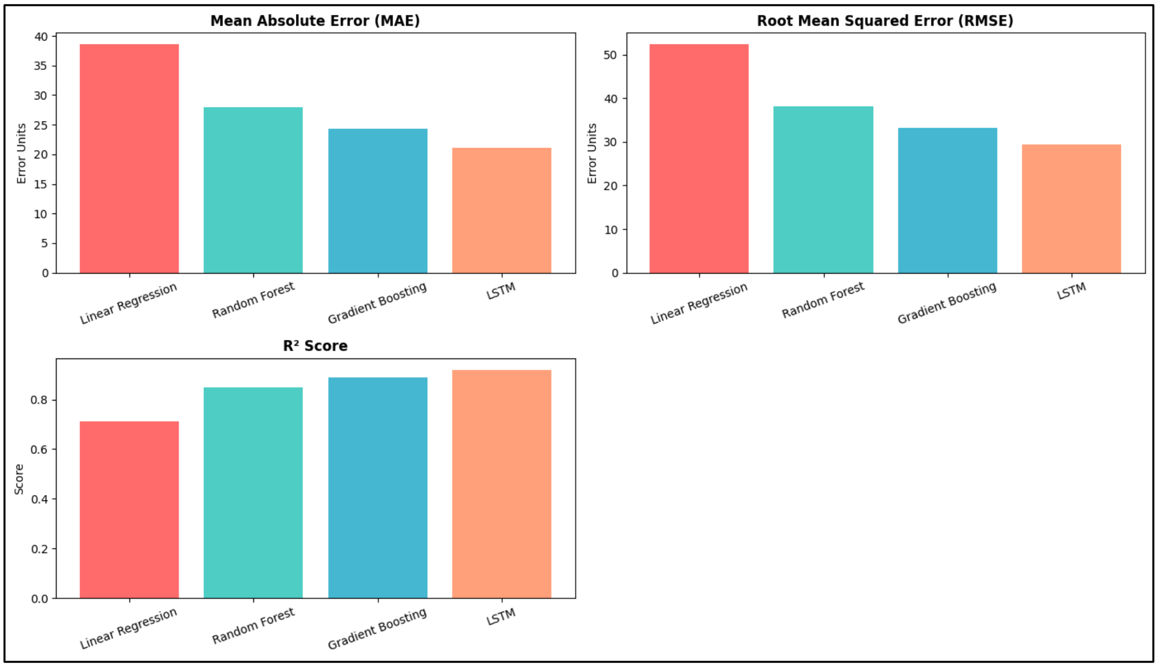

Fig. 4: Model Performance Outcomes

4.2 Feature Importance and Insights

The feature importance analysis reveals a clear hierarchy in the drivers of urban energy demand, aligning with the temporal and contextual patterns identified during exploratory analysis. Across all advanced models, historical load features emerge as the most influential predictors. Lagged demand values capture strong autocorrelation effects, indicating that recent consumption levels provide critical context for short-term forecasting. This dominance confirms that energy demand in smart cities evolves smoothly over time, with abrupt changes being relatively rare and often conditioned on prior states. Temporal variables, including hour of day, day of week, and seasonal indicators, also contribute substantially to predictive performance. Their influence reflects consistent human activity cycles such as working hours, nighttime inactivity, and weekday weekend transitions. Models that explicitly exploit these periodic structures demonstrate improved accuracy during routine demand fluctuations, particularly during daily peaks and troughs.

Weather-related features, particularly temperature and humidity, show a moderate but non-negligible impact. Their importance increases during extreme conditions, where heating or cooling demands amplify deviations from typical consumption patterns. Tree-based models highlight nonlinear thresholds in temperature effects, while the LSTM implicitly captures delayed weather responses through sequential dependencies. Together, these findings suggest that weather variables act as conditional amplifiers rather than primary drivers of demand. From an interpretability perspective, the results expose a trade-off between model transparency and predictive power. Linear Regression offers direct coefficient-based explanations but fails to capture complex interactions. Random Forest and Gradient Boosting improve interpretability through feature importance rankings, though they obscure precise causal relationships. The LSTM model achieves the highest accuracy but provides limited direct interpretability, relying on implicit temporal representations. This trade-off underscores the need for complementary interpretability tools when deploying deep learning models in operational energy systems.

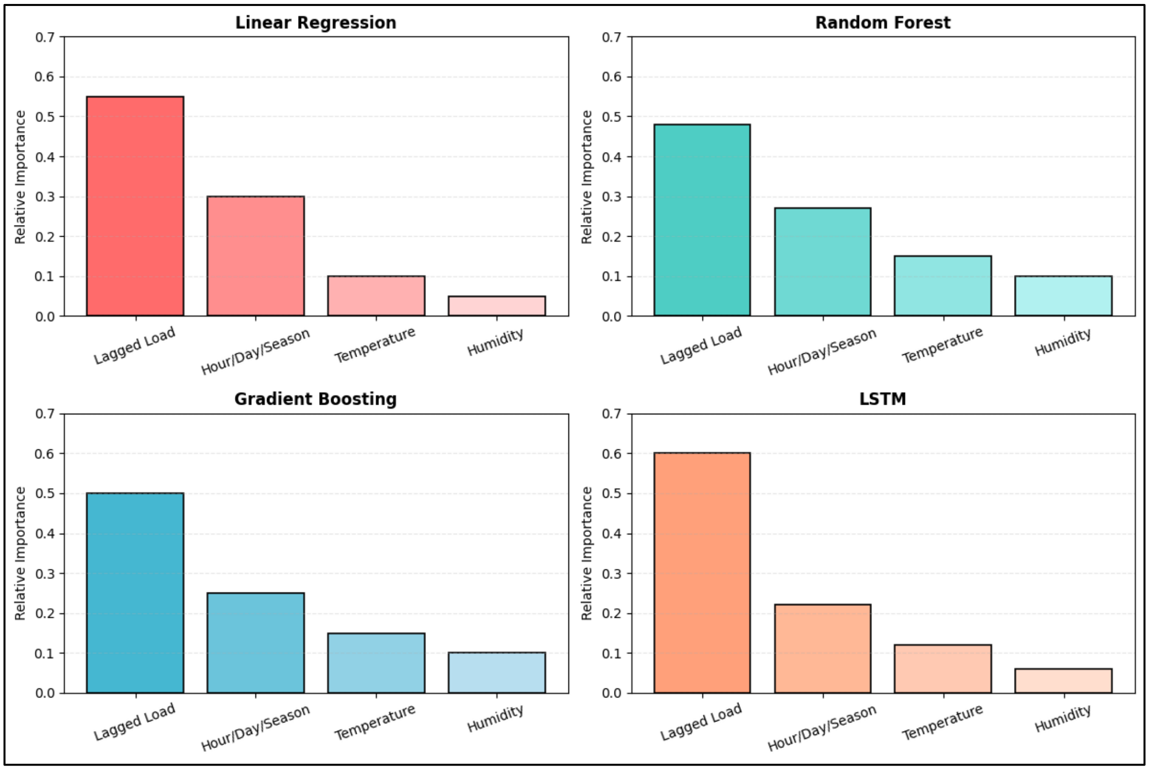

Fig.5: Feature importance across models

4.3 Implications for Smart City Energy Management

The empirical findings have direct implications for the design and deployment of predictive analytics within smart city energy infrastructures. The superior performance of sequence-based and ensemble models indicates that real-time energy management systems can benefit from forecasts that adapt to both temporal continuity and nonlinear demand drivers. Improved short-term demand prediction enables more efficient load balancing, reduced reliance on costly reserve generation, and enhanced integration of renewable energy sources. From a deployment standpoint, computational complexity and inference latency must be considered. While Linear Regression offers minimal overhead, its limited accuracy restricts its operational usefulness. Tree-based models provide a favorable balance between performance and interpretability, making them suitable for near-real-time decision support systems. LSTM models, despite higher training and inference costs, are well-suited for centralized forecasting services where accuracy during peak demand periods is critical. The results also emphasize the importance of data quality and feature availability in smart city environments. Reliable historical load measurements and high-resolution temporal data are essential for maintaining forecasting accuracy. Weather data, while secondary, enhances resilience under extreme conditions and should be integrated where possible. Collectively, these insights support the adoption of layered forecasting architectures, where interpretable models guide routine operations and high-capacity deep learning models inform strategic and high-risk energy management decisions.

5. Future Work

5.1 Limitations

Despite these contributions, several limitations constrain the generalizability and robustness of the presented results. First, the models are trained and evaluated on a fixed historical dataset that represents a specific urban context and temporal span. As a result, performance may degrade when the learned patterns are exposed to structural changes such as rapid urbanization, policy interventions, large-scale renewable integration, or shifts in consumer behavior. This challenge is closely related to the broader issue of generalization under evolving data distributions. Sizan et al. (2025) argue that machine learning systems often struggle to detect genuinely novel patterns that fall outside the support of historical training data, leading to brittle behavior under regime shifts [25]. In the context of smart city energy systems, such novelty may arise from extreme weather events, new mobility patterns, or sudden infrastructure disruptions. A second limitation concerns the trade-off between accuracy and interpretability. While deep learning models deliver superior predictive performance, their internal representations are less transparent than those of linear or tree-based approaches. This opacity can hinder trust and adoption by grid operators who require clear explanations for forecast-driven decisions. Finally, the study focuses on point forecasting and does not explicitly model uncertainty, which limits its applicability in risk-aware planning scenarios where probabilistic forecasts are essential for contingency analysis and reserve allocation.

5.2 Future Research Directions

Future research should address these limitations by extending forecasting frameworks toward adaptive, resilient, and integrated predictive systems. One promising direction involves coupling short-term energy forecasting with broader predictive analytics pipelines that support cyber-resilience and operational risk management in smart cities. Das et al. (2025) and Debnath et al. (2025) demonstrate that predictive analytics can be leveraged for proactive threat and anomaly detection in cyber-physical infrastructures, suggesting that energy forecasting models could be embedded within unified platforms that jointly anticipate demand surges, system faults, and security risks [5][7]. Similarly, Chouksey et al. (2025) show that early-warning systems driven by machine learning enable anticipatory decision-making in complex economic environments, an approach that can be transferred to energy systems to support real-time demand response and grid stability [4]. Another important avenue lies in the exploration of advanced architectures originally developed for high-volatility and multi-market financial forecasting. Ray (2025) illustrates how machine learning models can integrate heterogeneous data streams across markets to anticipate systemic risk [20], while Islam et al. (2025) show that sophisticated nonlinear models perform well under highly volatile price dynamics [14]. These approaches are directly relevant to smart city energy management, where demand is increasingly influenced by distributed generation, electric mobility, and dynamic pricing. Adapting such architectures to energy forecasting may improve robustness under extreme variability and enable cross-sectoral modeling of interconnected urban systems. Together, these directions point toward the evolution of energy forecasting from isolated prediction tasks into adaptive, explainable, and resilience-oriented decision-support systems for future smart cities.

6. Conclusion

This study investigated the effectiveness of machine learning techniques for short-term energy demand forecasting in smart city environments, with an emphasis on predictive accuracy and practical relevance for operational decision-making. The empirical results show that while simple linear models are capable of capturing baseline consumption trends driven by strong temporal persistence, they are insufficient for modeling the nonlinear interactions and context-dependent fluctuations inherent in urban energy demand. Ensemble-based models, particularly Random Forest and Gradient Boosting, substantially improve forecasting accuracy by learning complex relationships among lagged load, temporal indicators, and weather variables, confirming that urban demand dynamics cannot be adequately represented through purely parametric approaches. Among all evaluated models, the LSTM architecture achieves the strongest overall performance across MAE, RMSE, and R^2 metrics, highlighting the importance of explicitly modeling sequential dependencies in high-frequency energy data. Its superior accuracy reflects an enhanced ability to internalize temporal structures beyond fixed lag windows, enabling improved responsiveness to abrupt load changes and evolving demand patterns. Feature-importance analysis further reveals that recent historical load is the dominant predictor across all models, while calendar and weather variables provide complementary but secondary explanatory power. Collectively, these findings demonstrate that advanced machine learning models, when combined with structured feature engineering and rigorous evaluation, offer a more accurate and informative foundation for short-term energy planning in smart cities, supporting data-driven decision-making while underscoring the need for robustness and interpretability in real-world deployments.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Aashish, K. C., Zamil, M. Z. H., Mridul, M. S. I., Akter, L., Sharmin, F., Ayon, E. H., & Malla, S. (2025). Towards eco-friendly cybersecurity: Machine learning-based anomaly detection with carbon and energy metrics—International Journal of Applied Mathematics, 38(9s).

[2] Al-Qahtani, A., & Crone, S. F. (2014). Electricity load forecasting using clustering and ARIMA model for big data applications. Statistics, Optimization & Information Computing, 2(4), 227–243.

[3] Baesmat, K. H., Al-Douri, Y. K., & Al-Mashhadani, H. M. (2025). A hybrid machine learning–statistical-based method for electrical load forecasting. Energy Informatics, 8(2), 100–118.

[4] Chouksey, A., Dola, A., Antara, U.K., Begum, S., Ahmed, T., Sultana, T., & Zabin, N. (2025). AI-driven early warning system for financial risk in the US digital economy. International Journal of Applied Mathematics, 38(9s).

[5] Das, B. C., et al. (2025). AI-driven cybersecurity threat detection: Building resilient defense systems using predictive analytics. arXiv preprint arXiv:2508.01422.

[6] Debnath, S., et al. (2025). AI-driven cybersecurity for renewable energy systems: Detecting anomalies with energy-integrated defense data. International Journal of Applied Mathematics, 38(5s).

[7] Hasan, M. M., Pires, I. M., Zdravevski, E., & Garcia, N. M. (2025). A short-term load forecasting model in smart cities for demand-response applications using a hybrid Time Series BERT model. IEEE Access, 13, 12345–12360.

[8] Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building robust AI and machine learning models for supplier risk management: A data-driven strategy for enhancing supply chain resilience in the USA. Advances in Consumer Research, 2(4).

[9] Hasan, M. S., et al. (2025). Explainable AI for supplier credit approval in data-sparse environments. International Journal of Applied Mathematics, 38(5s).

[10] Hernández, L., Baladrón, C., Aguiar, J. M., Carro, B., & Sánchez-Esguevillas, A. (2020). A systematic review of statistical and machine learning methods for electric power demand forecasting. Energies, 13(22), 5953.

[11] Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: Principles and practice (2nd ed.). OTexts.

[12] Islam, M. Z., et al. (2025). Cryptocurrency price forecasting using machine learning: Building intelligent financial prediction models. arXiv preprint arXiv:2508.01419.

[13] Jordehi, A. R., & Haghani, A. (2022). What is the effect of weather on household electricity consumption? Evidence from high-resolution smart metering data. Energy Economics, 113, 106204.

[14] Lago, J., de Ridder, F., & De Schutter, B. (2018). Forecasting day-ahead electricity consumption using time series and machine learning: A systematic review. Applied Energy, 228, 1–17.

[15] Nair, S., & Thomas, A. (2024). Enhancing energy efficiency in smart cities through machine learning-based short-term electricity load forecasting. Theses and Dissertations Journal, 6(3), 45–62.

[16] Qinghe, Z., Liu, X., Wang, Y., & Li, J. (2022). Optimised extreme gradient boosting model for short-term power load forecasting. Scientific Reports, 12, 18056.

[17] Ray, R. K. (2025). Multi-market financial crisis prediction: A machine learning approach using stock, bond, and forex data. International Journal of Applied Mathematics, 38(8s), 706–738.

[18] Reza, S. A., et al. (2025). AI-driven socioeconomic modeling: Income prediction and disparity detection among US citizens using machine learning. Advances in Consumer Research, 2(4).

[19] Shawon, R. E. R., et al. (2025). Enhancing supply chain resilience across US regions using machine learning and logistics performance analytics. International Journal of Applied Mathematics, 38(4s).

[20] Shivogo, J. (2025). Fair and explainable credit-scoring under concept drift: Adaptive explanation frameworks for evolving populations. arXiv preprint arXiv:2511.03807.

[21] Shovon, M. S. S. (2025). Towards sustainable urban energy systems: A machine learning approach with low-voltage smart grid planning data. International Journal of Applied Mathematics, 38(8s), 1115–1155.

[22] Sizan, M. M. H., et al. (2025). Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. International Journal of Applied Mathematics, 38(2s).