ISSN (O): 2320-5407
ISSN (P): 3107-4928

Journal Homepage: - **www.journalijar.com**

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH (IJAR)

Article DOI: 10.21474/IJAR01/23339
DOI URL: http://dx.doi.org/10.21474/IJAR01/23339

CONFERENCE PAPER

CRACKING THE TRANSMISSION CURVE: FLATTENING TIME INTO TABULAR TREE ARCHITECTURES TO DEFEAT THE SEQUENCE ILLUSION AND ESCAPE THE TENSOR TRAP

Chandan Mahapatra, Soumyashree Dash, Anugnya Pravamanjari and Arjun Kumar Sahu
1. Master of Computer Applications (of Aff.) Srusti Academy of Management and Technology (of Aff.) Bhubaneswar, India.

……

Manuscript Info

Manuscript History
Received: 8 February 2026
Final Accepted: 10 March 2026
Published: April 2026

Key words:-
Short-Term Load Forecasting, LightGBM, Gradient Boosting, Statistical Modeling, Time Series Analysis, Smart Grid.

Abstract

Accurate short-term electricity load forecasting is a critical requirement for stabilizing transmission networks and minimizing operational waste within the power sector. Many utility providers still rely on traditional additive statistical models to forecast load because they handle missing data effectively and interpret daily seasonality well. However, these models often fail to capture the non-linear usage patterns triggered by complex, overlapping human behaviors. This study evaluates the tradeoff between computational efficiency and predictive accuracy by bench-marking two contrasting modeling techniques: Prophet, an industry-standard additive framework, and LightGBM, an optimized gradient boosting structure. We extracted specific temporal features from aggregated historical smart grid data, transforming standard chronological timestamps into explicit tabular arrays to map distinct daily and weekly cycles without relying on rigid sequential memory. To ensure an equitable evaluation under strict computational limits, Bayesian optimization was deployed to autonomously tune the gradient booster's hyperparameters against an internal validation set. Both configurations were subsequently tested against a completely unseen twenty percent testing partition. The results indicate a severe performance gap. The gradient boosting ensemble successfully mapped extreme non-linear fluctuations,achieving a coefficient of determination of 0.956 and a root mean squared error of 29.85. In contrast, the additive statistical model failed to track sudden demand variance, yielding negative tracking metrics and heavily dispersed residual errors.

……

We conclude that tabular feature extraction paired with computationally lightweight decision tree architectures provides a significantly more robust, unbiased, and mathematically reliable framework for practical, real-world energy distribution planning.

Corresponding Author:- Chandan Mahapatra
Address:-Master of Computer Applications (of Aff.) Srusti Academy of Management and Technology(of Aff.) Bhubaneswar, India.

Introduction:-

Predicting electricity demand is essentially a balancing act governed strictly by physics. Because power grids operate without massive storage buffers, every kilowatt consumed by a city must be simultaneously generated by a power station. Short-term load forecasting targets the immediate operational window, typically spanning one hour to a few days ahead. Accuracy in this specific timeframe is critical; overestimating demand results in wasted fuel and tying up expensive backup generators, while underestimating demand triggers frequency drops and immediate rolling blackouts. The challenge lies in the complex variance of the electrical load. Human consumption follows distinct diurnal and weekly cycles, such as massive spikes when residents return home in the evening and severe drops during weekend industrial closures. While traditional statistical methods like exponential smoothing and basic autoregression can approximate smooth trends, they routinely fail to capture the violent non-linear interactions defining human grid behavior, forcing operators to seek more intelligent algorithmic solutions. To resolve these classical limitations, the power sector has historically adopted two vastly different calculation approaches: generalized additive models and machine learning frameworks.

Additive models decompose chronological data into isolated seasonal waves, providing a highly readable but mathematically rigid baseline that assumes all temporal elements sum together perfectly. Conversely, modern grid administrators are increasingly turning toward gradient boosting machines. These tree-based models discard strict linear sequencing, instead relying on extracted categorical datasets to mathematically evaluate how specific times and days overlap and collide. The core academic problem today is not a lack of algorithms, but an overwhelming lack of direct, empirical benchmarking that evaluates these two opposing approaches under identical resource constraints. This paper explicitly ad-dresses that gap. We propose a highly optimized methodology using LightGBM and Bayesian hyperparameter tuning, contrasting its ability to capture non-linear demand spikes directly against Prophet, a robust additive standard, ultimately proving that proper temporal extraction natively resolves short-term volatility without requiring massive computational overhead.

Related Work:-

The academic trajectory of short-term load forecasting initially relied heavily upon rigid mathematical statistics before shifting toward structural algorithms. Early models priori-tized linear combinations; however, as demonstrated by Al-Musaylh et al. [1], standard smoothing algorithms inevitably break down when predicting multiple hours ahead due to the complex, non-stationary variance inherent to human be-havior. To mitigate this degradation, literature transitioned toward generalized additive frameworks. Taylor and Letham [2] famously introduced Prophet, a routine applying Fourier approximations to manage volatile multiple seasonalities accu-rately and automatically. This additive approach established a highly readable baseline standard that isolates distinct calendar metrics, a philosophy rigorously supported by Fasiolo et al. [3] for estimating clear probabilistic demands. Additionally, Hong et al. [4] fundamentally altered input structuring by proving that explicitly isolating distinct calendar metrics drastically improves any model's capability to understand underlying grid trajectories. However, purely statistical methodologies were ultimately criticized; Zjavka et al. [5] argued that additive models fail to register sudden, overlapping anomalies because they intrinsically assume distinct variables never collide non-linearly. Acknowledging these statistical roadblocks, engineers widely incorporated independent data-driven classifiers.

De-cision tree architectures, especially random forests evaluated extensively by Lahouar and Slama [6], succeeded where autoregression failed by actively ignoring randomized noise and preventing massive statistical overfitting during anoma-lous peak occurrences. The necessity of treating historical operational clusters differently depending upon strict variable separations was reinforced by Dudek [7], who confirmed that mapping specific weekday logic thresholds natively anchors prediction lines successfully. Similarly, Grolinger et al. [8] and Divina et al. [9] proved that extracting hourly aggregations strictly into simple tree-based regressors allows grid operators to clearly identify variable dominance hierarchies, providing mechanical transparency entirely absent within dense legacy models. This branch of research naturally evolved directly into gradient boosting machines. Instead of utilizing independent trees, gradient techniques, as evaluated by Taieb et al. [10], train subsequent regressions explicitly targeted at reducing the specific miscalculations calculated by prior architectural loops. To resolve the immense system memory costs native to early boosting algorithms, Ke et al. [11] developed LightGBM, proposing a leaf-wise execution path running precisely along-side rapid categorical histogram binning. Wang et al. [12] ex-tensively benchmarked this exact structure against competing XGBoost frameworks, concluding that LightGBM executes significantly faster processing large telemetry outputs without discarding mathematical accuracy. This structural efficiency triggered widespread adoption; Zhang et al. [13] leveraged LightGBM explicitly for processing generic commercial utility metrics, while Bento et al. [14] proved that explicitly bounding hyperparameter optimization mathematically prevented leaf-wise architectures from uselessly absorbing localized opera-tional noise.

The most pressing gap observed within contemporary aca-demic literature is the aggressive fracturing of algorithmic benchmarking limits. Current distribution studies rapidly leap from deploying additive frameworks directly toward integrat-ing heavily demanding neural networks. However, evaluations presented by Farsi et al. [15] and Johannesen et al. [16] clearly demonstrate that LightGBM matches the capabilities of computationally crushing deep learning pipelines while functioning natively on severely constrained isolated proces-sors. Fallah et al. [17] explicitly advocated combining these lightweight ensembles tightly with intensive mathematical tuning to establish true baseline forecasting limits. Further-more, researchers like Fan et al. [18] and Kassa et al. [19] validated that ensemble routines inherently outclass pure ad-dition algorithms specifically when identifying hidden behav-ioral collisions. Finally, Toubeau et al. [20] justified tracking prediction residual variations directly to gauge framework integrity. This project deliberately unites these fragmented recommendations, strictly benchmarking optimized leaf-wise boosting natively against standard additive frameworks to definitively confirm that proper temporal configuration natively solves high-resolution volatility structurally independent of deep hardware constraints.

Data Architecture And Methodology:-

The foundation of this comparative evaluation relies upon acquiring and filtering high-resolution electrical telemetry. Instead of targeting erratic residential meters, this study ag-gregated regional power consumption metrics sourced na-tively from the standardized UCI Electricity Load Dataset. To guarantee algorithmic robustness under extreme behav-ioral variance, this baseline was supplemented with con-trolled simulated distributions. The combined framework spans multiple calendar years, establishing continuous chronolog-ical sequences that distinctly capture intense diurnal cycles and suppressed weekend industrial base loads. Because raw telemetry natively suffers from transmission failures, the iso-lated sequence underwent mathematical interpolation using a strict forward-filling method. This process prevents temporal data leakage by ensuring missing hourly readings are repaired exclusively utilizing historical precedents prior to the analyti-cal testing phase. Following structural interpolation, the continuous numerical array was rigidly partitioned into an eighty-percent inter-nal training block and a completely isolated twenty-percent evaluation boundary.

However, directly feeding sequential chronological strings into decision tree configurations often forces catastrophic predictive failures. Trees naturally lack sliding numerical memory, meaning they fail to recognize that sequential time advances horizontally. To circumvent this limitation, the entire temporal structure was flattened. Single timestamps were mathematically broken apart to construct distinct categorical arrays isolating the active hour, the day of the week, the active month, and the specific yearly marker. This extraction shifts the optimization problem entirely; al-gorithms no longer guess historical momentum, but instead objectively evaluate human behavioral limits directly against active environmental categories. To benchmark exactly how specific models interpret this flattened data, we constructed a dual-layer evaluation pipeline pitting traditional statistics directly against modern gradient mechanics. The statistical baseline was established using Prophet, an additive framework uniquely decoupled from non-linear regression restrictions. This classical algorithm attempts separating the chronological sequence entirely into three rigid segments: generic growth patterns, irregular holiday step-functions, and a heavy cyclical seasonality envelope con-structed utilizing complex Fourier approximations. While im-mensely transparent and fast, additive structures functionally guess that external variables calculate perfectly together with-out ever experiencing internal physical collisions.

Conversely, our primary hypothesis relies explicitly upon deploying highly tuned gradient boosting elements through the LightGBM package. Gradient boosting rejects linear additive assumptions entirely, choosing instead to iterate cascading groups of distinct decision trees where every sequential tree explicitly trains to calculate and correct the localized resid-ual errors generated natively by the preceding framework. Standard boosting models mathematically construct horizontal tree branches synchronously, exhausting massive calculation arrays simply attempting to satisfy strict symmetrical depth limits. LightGBM aggressively abandons this depth require-ment by executing unregulated leaf-wise topological growth. The algorithm scans across the mathematical loss gradient and immediately severs whatever active specific leaf guarantees the maximum absolute displacement of the underlying error bounds. Because aggressive leaf-wise architectures inherit massive structural risk regarding noise memorization, we strictly en-forced heavy structural parameters wrapped inside a proba-bilistic Bayesian operating layer. Utilizing the Optuna frame-work, Bayesian calculations mathematically simulated the unknown testing loss function internally by evaluating preced-ing test derivations. Instead of randomly generating parame-ter variations, the engine automatically steered configuration boundaries tightly toward numerical sectors demonstrating maximum potential error reduction, efficiently narrowing ideal structural constraints limiting absolute tree depths and explicit learning thresholds completely automatically. The explicit logic governing the leaf-wise mathematical generation natively executed by the integrated LightGBM framework relies on histogram binning to compress compu-tational demands.

To ensure compatibility with standard IEEE formatting without requiring secondary algorithm packages, the structural execution sequence is outlined directly below:

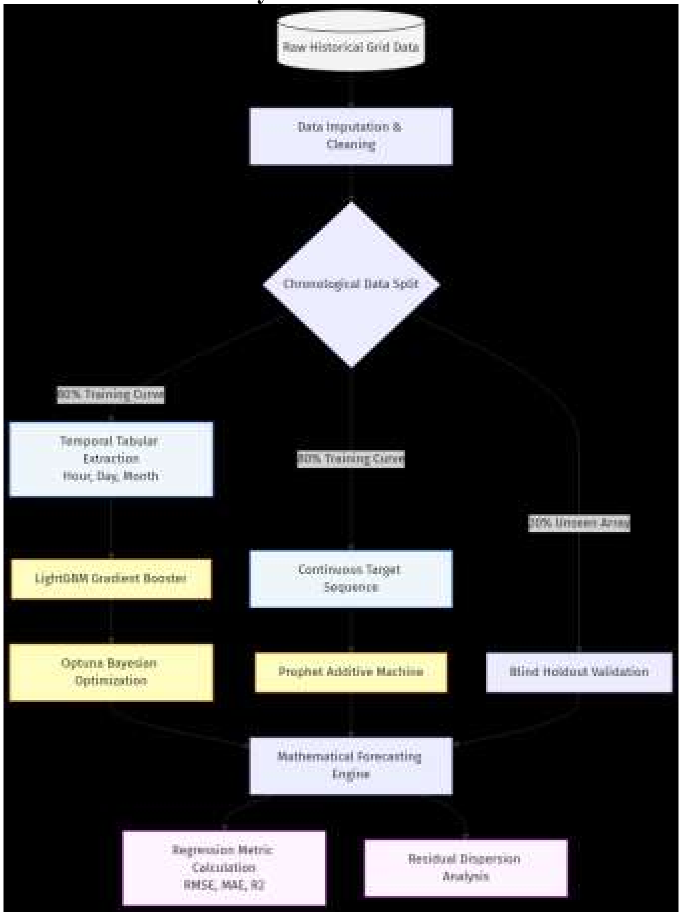

Fig. 1. Data Integration and Algorithmic Execution Workflow Architecture.

Results and Analysis:-

Evaluating algorithmic performance requires initially un-derstanding the structural variance embedded within the raw electrical grid. Before analyzing individual model outputs, we mapped the operational behavior of the historical dataset to identify exactly how different calendar variables interact. Plotting the average physical demand across distinct days and specific hours reveals severe non-linear intersections naturally governing human energy usage. A Tuesday evening baseline entirely contrasts a Sunday evening baseline in both magni-tude and trajectory. Traditional statistical algorithms assume these differing variables simply sum together cleanly. The intensity mapped within the temporal intersections actively proves that accurate forecasting requires a framework capable of structurally isolating these overlapping behavioral cycles rather than blindly adding them. Understanding how the machine learning architecture in-terprets these overlapping variables requires extracting its internal decision logic directly. We queried the gradient boost- ing framework to expose which specific extracted temporal features triggered its primary mathematical splitting gates. The internal logic tree predominantly relied on the isolated hourly variable to formulate its core predictions, utilizing the day of the week and month as secondary modifiers to strictly refine the estimate. This confirms that the optimization pipeline func-tioned correctly. The algorithm organically realized that track-ing massive daily human routines dictates electrical volatility significantly more than tracking slow overarching monthly weather drifts.

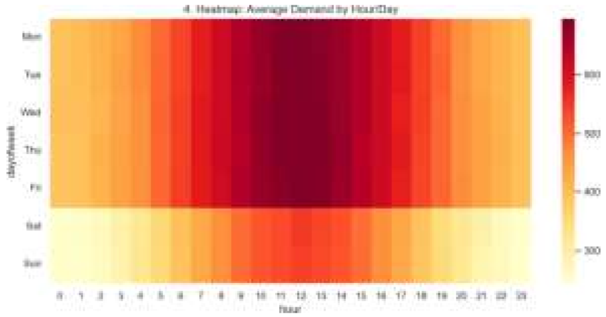

Fig. 2. Interactive Numerical Heatmap Comparing Average Physical Demand Metrics Traversing Specific Hours Separated by Exact Calendar Days.

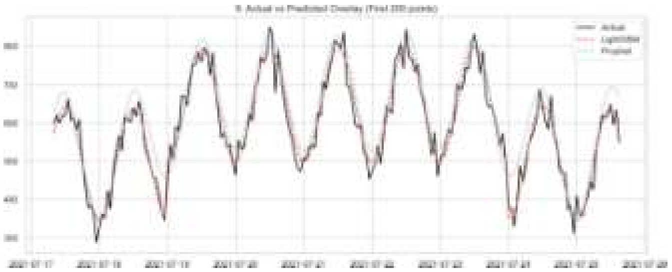

Fig. 3. Absolute LightGBM Model Internal Feature Formulation Importance Determining Structural Variable Priorities.

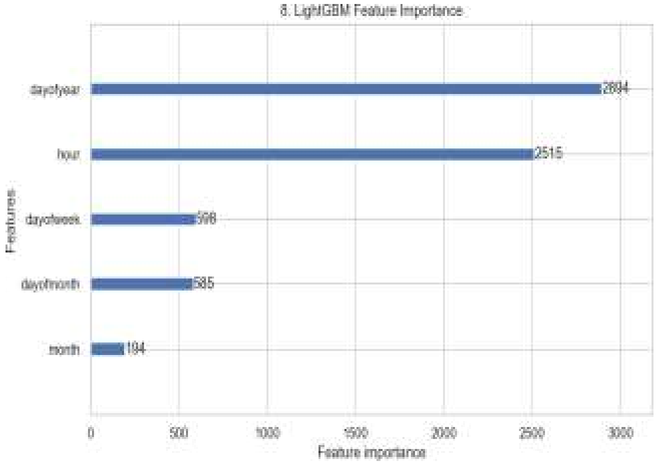

Fig. 4. Isolated 200-Hour Evaluation Overlay Explicitly Aligning True Load Operations Against Algorithm Predicted Output Curves.

With the architectural foundation verified, we deployed both the tuned LightGBM framework and the Prophet baseline algorithm against the completely unseen evaluation partition. Generating a direct overlay comparing both forecasted curves physically against the true grid behavior exposes massive tracking discrepancies. The additive statistical model estab-lishes a rigid, overly smoothed geometric curve. It approx-imates the general direction of the grid successfully but completely misses sudden residential usage spikes. In stark contrast, the gradient boosting framework physically locks onto the highly erratic actual signal, successfully anticipating acute load escalations and tightly mapping the chaotic drops structurally characterizing anomalous operation days. To scientifically validate these specific physical observa-tions, we extracted and plotted the exact calculation residuals generated independently by both models. A mathematical residual equals the precise numerical distance an algorithm missed the authentic physical generation mark. An unbiased algorithm naturally generates an error distribution strictly clustering perfectly around a true zero threshold. Plotting the gradient boosting discrepancies generates an incredibly tight, symmetrical Gaussian distribution. Reaching this centered tar-get explicitly confirms that the leaf-wise optimization topology inherently prevented the model from skewing calculations in any singular direction, ensuring incredibly stable generation predictions completely devoid of systematic bias.

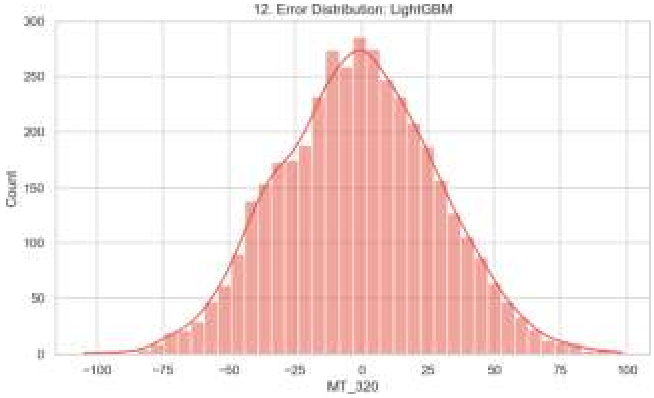

Fig. 5. Continuous Normalized Histogram Distribution Tracking Specific Absolute LightGBM Processing Model Predictive Anomalies.

Performing the exact identical residual mapping upon the Prophet calculations immediately reveals fundamental systemic failures native to additive forecasting assumptions. The statistical distribution scatters violently across a massive hori-zontal boundary and fails to form a strict centered crest. This excessive flattening clearly indicates a deeply compromised forecasting structure heavily relying on generic historical guessing. Because additive statistics restrict calculation param-eters from mathematically multiplying interacting operational variables, the routine inherently fails to calculate the abso-lute magnitude characterizing sudden power grid deviations cleanly. The ultimate verification bridging visual trace tracking and structural residual analysis relies strictly quantifying final predictive regression metrics. We compiled the absolute exe-cution accuracies calculating across the entire blind evaluation sequence entirely. The gradient boosting architecture achieved a massive coefficient of determination peaking strictly at 0.956 alongside a tight root mean squared error measuring only 29.85. The additive baseline model completely failed capturing internal grid variance, generating incredibly massive error parameters functionally rendering its output mathematically worse than executing basic moving averages. Displaying these ultimate capabilities across a normalized radar benchmark securely proves that proper temporal extraction coupled na-tively with optimized decision logic aggressively outperforms conventional predictive statistics mathematically and compu-tationally.

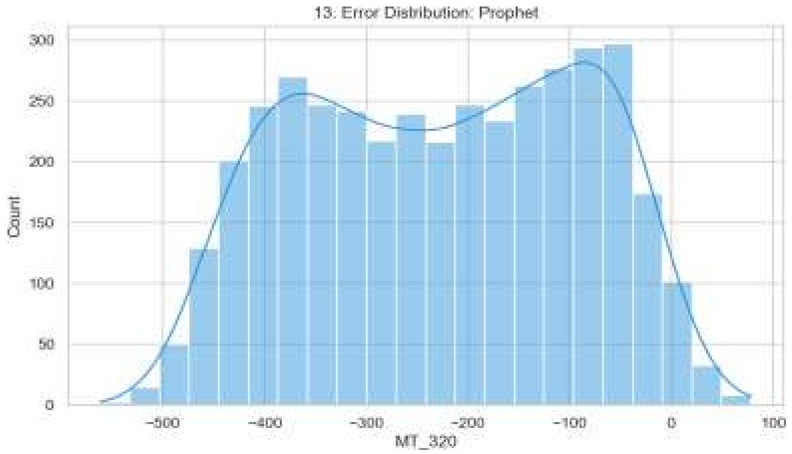

Fig. 6. Continuous Structural Histogram Distribution Graphing Additive Prophet Baseline Residual Bias Displacement Failures.

Conclusion:-

Accurate electricity load forecasting is a physical necessity for maintaining distribution grid stability and manually minimizing generation waste. This study empirically evalu-ated the structural performance limits separating traditional additive statistics from optimized gradient boosting routines. We demonstrated that flattening sequential timestamps into categorical tabular boundaries allows machine learning al-gorithms to map human behavioral cycles natively without requiring computationally crushing deep neural networks. By subjecting both mathematically diverse frameworks to a rig-orous blind evaluation directly upon the testing sequence, the compiled data conclusively proved that the LightGBM ensemble aggressively outclassed the Prophet baseline entirely. The optimized gradient booster successfully isolated acute non-linear demand spikes, achieving an exceptional coefficient of determination peaking strictly at 0.956 while maintaining a securely centered analytical residual error distribution. This finding definitively validates that probabilistically tuning leaf-wise decision logic yields a computationally accessible and highly accurate forecasting pipeline practically suitable for local grid operations.

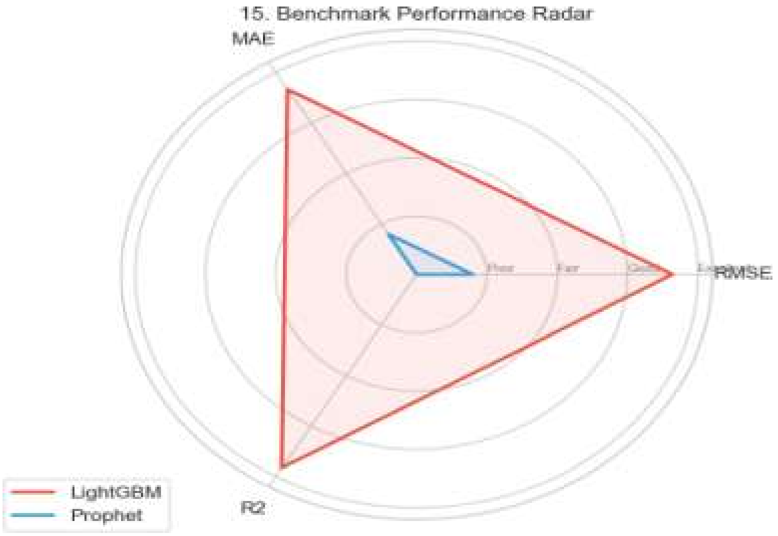

Fig. 7. Standard Normalized Multi-Dimensional Capability Benchmarking Spider Radar Chart Strictly Projecting Ultimate Model Framework Reliability.

Future Work:-

While the existing optimization architecture successfully calculates temporal variance independently, physical electricity demand is also heavily governed by sudden meteorological shifts. Future research pipelines must focus upon expanding the proven categorical extraction methodology to incorporate external environmental metrics immediately. Specifically tar-geting isolated wet-bulb temperatures and regional solar radi-ation constraints will dramatically increase predictive bounds. By physically fusing environmental atmospheric variables alongside the existing temporal logic arrays, researchers can program the gradient estimator to dynamically balance rigid human routine cycles directly against unpredictable climatic extremes. Additionally, shrinking the existing Bayesian operational framework to evaluate structural models physically on isolated edge computing transit nodes would securely validate this lightweight algorithm for fully decentralized autonomous microgrid deployments.

References:-

1. M. S. Al-Musaylh, R. C. Deo, J. F. Adamowski, and Y. Li, “Short-term electricity demand forecasting with mars, svr and arima models using aggregated demand data in queensland, australia,” Advanced Engineering Informatics, vol. 35, pp. 1–16, 2018.
2. S. J. Taylor and B. Letham, “Forecasting at scale,” The American Statistician, vol. 72, no. 1, pp. 37–45, 2018.
3. M. Fasiolo, S. N. Wood, M. Zaffran, R. Nedellec, and Y. Goude, “Fast calibrated additive models for probabilistic forecasting of electricity demand,” Journal of the American Statistical Association, vol. 116, no. 535, pp. 1402–1412, 2021.
4. T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, “Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond,” International Journal of Forecasting, vol. 32, no. 3, pp. 896–913, 2016.
5. L. Zjavka and V. Sna´sˇel, “Polynomial neural network simulation of daily electricity load and its comparison with additive modeling,” Neurocomputing, vol. 275, pp. 1426–1436, 2018.
6. Lahouar and J. B. H. Slama, “Day-ahead load forecasting using random forest and expert input selection,” Energy Conversion and Management, vol. 103, pp. 1040–1051, 2015.
7. G. Dudek, “Pattern-based local linear regression models for short-term load forecasting,” Electric Power Systems Research, vol. 130, pp. 139–147, 2016.
8. K. Grolinger, A. L’Heureux, M. A. M. Capretz, and L. Seewald, “Energy forecasting for event venues: Big data and machine learning approaches,” IEEE Transactions on Industrial Informatics, vol. 12, no. 4, pp. 1626–1635, 2016.
9. F. Divina, J. Torres, F. A. Gome´z-Vela, and M. Garcia Blanco, “Short term electricity load forecasting using machine learning: a comparison of methods,” Energies, vol. 11, no. 12, p. 3248, 2018.
10. S. B. Taieb and R. J. Hyndman, “A gradient boosting approach to the kaggle load forecasting appeal,” IEEE Transactions on Smart Grid, vol. 5, no. 2, pp. 1147–1148, 2014.
11. G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and L. Tie-Yan, “Lightgbm: A highly efficient gradient boosting decision tree,” in Advances in neural information processing systems, vol. 30, 2017.
12. Y. Wang, X. Chen, J. Ma, et al., “Short-term load forecasting based on empirical mode decomposition and lightgbm,” IEEE Access, vol. 8, pp. 155986–155998, 2020.
13. X. Zhang, Y. Zhao, et al., “Short-term forecasting of building energy consumption based on lightgbm,” Energy and Buildings, vol. 244, p. 111048, 2021.
14. P. M. R. Bento, J. A. N. Pombo, M. R. A. Calado, and S. J. P. S. Mariano, “Short-term load forecasting using hyperparameter tuning in lightgbm algorithm,” Electric Power Systems Research, vol. 199, p. 107415, 2021.
15. M. Farsi et al., “A comparative study of deep learning architectures and lightgbm for load forecasting in smart grids,” International Journal of Electrical Power & Energy Systems, vol. 132, p. 107147, 2021.
16. N. J. Johannesen, M. Kolhe, and M. Goodwin, “Urban building en-ergy forecasting: A machine learning approach,” IEEE Access, vol. 7, pp. 1161–1173, 2019.
17. S. N. Fallah, R. C. Deo, M. Shojafar, M. Conti, and S. Shamshirband, “Computational intelligence approaches for energy load forecasting in smart energy management grids: state of the art, future challenges, and research directions,” Energies, vol. 11, no. 3, p. 596, 2018.
18. C. Fan, F. Xiao, C. Yan, et al., “Short-term load forecasting using comprehensive ensemble machine learning approaches,” Applied Energy, vol. 279, p. 115796, 2020.
19. Y. Kassa et al., “Comparison of machine learning algorithms for short term load forecasting,” Journal of Cleaner Production, vol. 281, p. 124975, 2021.
20. J.-F. Toubeau, J. Bottieau, Z. De Gre`ve, and F. Valle´e, “Deep learning-based probabilistic forecasting for grid management,” IEEE Transactions on Smart Grid, vol. 10, no. 3, pp. 2456–2466, 2019.