Hybrid Gradient Boosting and Anomaly Detection Framework for Autonomous Peak Demand Management in Commercial Building Energy Systems

Pruthvi Raj Seknametla
Independent Researcher
pruthviraj9369@gmail.com

Hari Suresh Babu Gummadi
Independent Researcher
Texas, USA
hari5950@gmail.com

Sankalp M R
Assistant Professor, CSE-ICB
Dayananda Sagar Academy of Technology and Management, Bangalore, Karnataka
sankalp-icb@dsatm.edu.in

Mugunth Chandran
Independent Researcher
mugunth.visa@gmail.com

A Mutharasan
Associate Professor
Department of Electronics and Communication Engineering
Vel Tech Rangarajan Dr Sagunthala R&D Institute of Science and Technology
Chennai, Tamil Nadu, India
amutharasan@veltech.edu.in

Sushma Sunkollu Nagaraj
Solutions Architect, USA
snsushma94@gmail.com

***Abstract*— Forecast and anomaly detection are two modules that most building energy management systems run separately, this paper tests what happens when they joined together. We comparatively assessed an integrated framework comprising Light. GBM consumption forecasting with Isolation Forest peak detection on 8760 hourly observations from a research facility in Porto, Portugal, in which five decision integration strategies were evaluated on an identical temporal test set. The weighted mix (0.6×forecast score + 0.4×anomaly score) scored F1 = 0.9028 — an improvement of 8.1% over standalone anomaly detection (F1 = 0.822) with precision = 98.5%, and recall =84.3%. The forecasting module alone achieved MAE = 55.37 kW. The most unexpected result was the moderate error (50–200 kW) interaction with forecast error, which exhibited improved detection performance over low-error hours based on F1 score (F1 = 0.868 vs. F1 = 0.725): indicating that these two modules capture different and complementary signals of consumption. Economic evaluation predicted annual savings of €137,934 and system costs of €5,000 per year resulting in a 9.6:1 benefit-cost ratio. Inference latency of <500ms indicates that the framework has become real-time autonomous and can now be deployed on standard building automation hardware.**

I. INTRODUCTION

1.1 Motivation and Background

Demand charge penalties (calculated on peak usage over a billing cycle rather than total energy consumed) can constitute 30–70% of a commercial building's monthly electricity bill [1]. Getting these peaks before they happen, or at least the moment they start will certainly be worth a lot[2]. The most trouble is that it treats two totally different parts as completely separate modules: one for forecasting and planning in advance; another for threshold alarm and real-time action [3]. Having these two things separated is a waste. A prediction that knowledge upsurge in consumption would be coming just when the anomaly detector is most sensitive[4]. An anomaly signal, on the other hand, raises a flag that there has been some change in conditions sufficient to necessitate forecast re-calibration[5]. But does blending these two abilities into a single decision framework lead to measurably superior peak detection than either module when deployed independently — and if so, which integration strategy is most effective in practice? Our answers to both questions are supported with a real annual building dataset and full economic validation of multiple Mix-Use Energy (MUE) systems [6-7].

1.2 Literature Review

Gradient boosting methods—especially LightGBM and XGBoost for short-term building energy forecasting—are now the workhorse with non-linear temporal relationships being effectively managed using limited data, making hourly-ahead MAPE <2% [8]. They are computationally efficient and robust given a small amount of training data, thus suitable for operational deployment on building automation hardware where delay due to deep learning inference is problematic[9-10].

In fact, since October 2023, in practical applications for building energy anomaly detection Isolation Forest consistently outperformed density-based alternatives: no distance calculations slow it down as they do LOF and k-nearest neighbours methods making it perfectly scale to high-dimensional sensor data [11-12]. Solís-Villarreal et al. Confirmed Isolation Forest precision greater than 80% for energy outlier detection on properly tuned contamination thresholds specific to smart city deployments [13]. In demand response, Rizvi et al. Forecast-guided peak shaving, which can reduce demand by 10–30%, depending on building thermal mass and load flexibility [14], has so far been implemented with instantaneous threshold alarms, not predictive signals. El Gohary et al. In the context of economics, establishing recall as the right operational metric is important: a single missed peak will erase gains that are made by correctly identifying ten others — this fact was illustrated in [14-16].

This is a simple gap that this paper tackles, as prior works verify forecasting accuracy and anomaly detection F1-scores separately [17]. Going back to MPFR, nobody has systematically tested what overflowing these metrics gives you, which integration rule works best or whether moderate (as opposed to small) forecast errors could actually help

detection performance rather than detract from it [18-19]. The goal of this work is to fill each of these gaps with the empirical evidence and an overall economic validation.

1.3 Main Contributions

Five contributions this paper provides to the construction of building energy peak demand management:

Integrated system architecture: Integrated LightGBM and Isolation Forest framework simultaneously estimating peak demand in order to retrieve both predictive and reactive identification of peaks within a shared operational building energy managerial system to the best of our knowledge.

Methods of integrating decisions: The effect of five different integration strategies (anomaly-only, guided forecast AND decision-rule, either-or OR, weighted combination and a conservative threshold) are compared against each other. Forecast + Anomaly Weighted Combination wins with F1 = 0.9028, precision 98.5%, recall 83.3% — +8.1% gain over dedicated detection.

Forecast-detection interaction analysis: The result that moderate forecast errors (50–200 kW) yield better detection (F1 = 0.868) than little-error hours (F1 = 0.725) reveals the anti-intuitive relationship of forecasting and detection quality; often considered to be coupled modally, our results suggest these metrics do not co-vary positively. The two modules seem to be reflecting a variety of consumption signals.

Efficiency Assessment: Cost-benefit analysis calculates annual savings of €137,934 corresponding to a benefit-cost ratio of 9.6:1 — quantifying operational worth beyond F1-scores and validating commercial deployment.

Assessment of operational readiness: Sub-second (<500ms) inference latency, stable thresholds across seasonal transitions and a false alarm rate of 1.5 per 100 hours demonstrated that the framework was viable for industry deployment as an autonomous building automation solution without specialist hardware.

II. METHODOLOGY

2.1 System Architecture

We combine a LightGBM forecasting module with an Isolation Forest anomaly detector and integrate their outputs in five different decision integration strategies to create warnings for peak demand. The data is chronologically split — 70% training (6,132 hours), 15% validation (1,314 hours) and 15% test (1,314 hours)—and not shuffled at any stage. Feature engineering outputs 27 temporal predictors, such as lag features (1h, 24h, 168h consumption), rolling statistics (mean and standard deviation for the past day), and temporal encodings (hour of the day, day-of-week hour, month). All preprocessing parameters are fit only on training data.

2.2 Forecasting Module

Lightgbm regressor with 100 gradient-boosted trees, max depth 7, learning rate 0.1, and 31 leaf nodes per tree for 1-hour ahead consumption forecasts. With tree growth leaf wise and mean squared error loss, each successive tree corrects residual errors from the previous one. It can be deployed to the standard CPU hardware with < 1s inference latency.

2.3 Anomaly Detection Module

Isolation Forest; 100 estimators; 10% contamination = consumption patterns that stand out from operating normal The more random partitions needed to isolate an observation, the happier it is Normal observations (with shorter average path lengths) get lower anomaly scores. Before scoring, all 27 features are standardized to zero mean and unit variance using training-set statistics. Correspondingly, the observed maximum event frequencies along the dataset can be measured quantifying that, up to a 10% of contamination is obedient how much transgression takes over.

2.4 Decision Integration Strategies

The following five methods aggregate the forecasts and anomalous outputs into a single alert signal:

Anomaly-only: directly using Isolation Forest predictions as baseline.

Forecast guided AND: Fire alert only when forecast $\geqslant$ 90th percentile + anomaly score $\geqslant$ 90th percentile.

This OR fire alert condition is when either of the conditions fires.

Weighted combination: S_combined = 0.6×S_forecast + 0.4×S_anomaly, with both scores min-max normalized to [0,1] prior to fusion Weights learned from validation set by disabling subjects with minimum recall of 80% over F1 maximization using grid search.

Conservative: Performing 95th percentile thresholds on both components independently (best precision, worse recall).

2.5 Economic Evaluation Framework

The economic model addresses demand response across the simulation period. True positive detections cause a peak load curtailment of 10% (855.8 kW) since maximum consumption decreases from 8,558 kW to 7,702 kW only leading to demand charges of €15/kW/month. For every false positive event occurring, €50 is incurred in operational penalties namely; thermal comfort disruption, equipment cycling and occupant disturbance. Net savings are simply the demand charge reduction minus the cost of false positives. Annual estimates extrapolate performance in the test period across 12 bills in the billing cycle.

2.6 Implementation Details

Standard CPU hardware completes full training within 0.8 seconds (LightGBM 0.6s, Isolation Forest 0.2s). Inference on test data — Of course, this all happened as part of the process which needs to be able run real time results during a building automation / smart home automation use case (test inference runs at 80 Michroseconds per prediction!) Decision thresholds are calibrated by peforming grid search with the size of 5 percentiles on the validation set based on the range of 75th–99th percentiles.

III. MATHEMATICAL MODELING

3.1 LightGBM Forecasting

The forecasting module builds an additive ensemble of

M = 100 gradient-boosted trees:

$$\hat{y}_{t+h} = \sum_{m=1}^{M} \eta \cdot f_m(\mathbf{x}_t) \quad (1)$$

where $\eta = 0.1$ is the learning rate and f_m is the m-th tree's prediction. Training minimizes a regularized objective:

$$\mathcal{L} = \sum_{i=1}^{N} L\,(y_i, \hat{y}_i) + \sum_{m=1}^{M} \Omega\,(f_m) \quad (2)$$

where l is mean squared error and the regularization term $\Omega(f_m) = \gamma T_m + \frac{1}{2}\lambda \Sigma_j w_{m,j}^2$ penalizes tree complexity through leaf count T_m and leaf weights $w_{m,j}$.

3.2 Isolation-Based Anomaly Detection

Anomaly likelihood is quantified through the average path length needed to isolate each observation across 100 trees:

$$s(\mathbf{x}) = 2^{-\frac{E[h(\mathbf{x})]}{c(n)}} \quad (3)$$

where E[h(x)] is the average isolation path length and c(n) = $2H(n-1)-2(n-1)/n$ normalizes scores across datasets, with H denoting harmonic numbers. A score that is close to 1.0 refers to an anomaly whilst a score of ~0.5 means normal observations. The decision threshold θ_s is set to the 90th percentile of the training anomaly scores, conditioned on contamination parameter $\gamma = 0.10$.

3.3 Feature Engineering and Temporal Dependencies

The 27-dimensional feature vector at each hour combines lag consumption, rolling statistics, and temporal encodings:

$$\mathbf{x}_t = \left[y_{t-1}, y_{t-24}, y_{t-168}, \mu_{t,24}, \sigma_{t,24}, h_t, d_t, m_t\right] \quad (4)$$

where - y_{t-k} - is consumption k hours in advance, $\mu\{t,24, \}$ and $\sigma\{t,24\}$ are the hourly rolling mean and standard deviation respectively. h_t d_t m_t are hour day of week and month Before Isolation Forest scoring, both the training-set mean — and standard deviation are Standardized..

3.4 Decision Integration Strategies

The weighted combination fuses min-max normalized forecast and anomaly scores:

$$S_{\text{combined}}(\mathbf{x}_t) = w_f \cdot S_f(\hat{y}_{t+h}) + w_a \cdot s(\mathbf{x}_t) \quad (5)$$

where -

$$S_f(\hat{y}_{t+h}) = \frac{\hat{y}_{t+h} - \min(\hat{\mathbf{y}})}{\max(\hat{\mathbf{y}}) - \min(\hat{\mathbf{y}})} \quad (6)$$

normalizes forecast predictions to [0,1]. Weights were optimized using grid search maximising F1 (subject to recall $\geqslant 0.80$) on the validation set. An alert fires if A also fyres, or the combined score exceeds θ_p in the 90th percentile:

$$\text{Alert}(t) = \begin{cases} 1 & \text{if } S_{\text{combined}}(\mathbf{x}_t) \geq \theta_p \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Peak hours are defined as consumption at or above the 90th percentile:

Peak(t) = 1 if $y_t \geq P_th$ = 4,296.6 kW, else 0

3.5 Economic Cost-Benefit Formulation

Baseline demand charge over the test period (1.8 months):

$$C_{\text{baseline}} = P_{\text{max,baseline}} \cdot r_d \cdot n_{\text{months}} \quad (8)$$

True positive detections trigger 10% curtailment: $P_\{reduced,i\} = P_i \cdot (1 - \rho)$, $\rho = 0.10$, reducing the billing peak from 8,558 kW to 7,702 kW. Revised demand charge:

$$C_{\text{intervention}} = P_{\text{max,reduced}} \cdot r_d \cdot n_{\text{months}} \quad (9)$$

False positive events incur operational penalties [10]:

$$C_{\text{FP}} = FP \cdot C_{\text{FP,unit}} \quad (10)$$

The net economic benefit combines gross savings minus penalties:

$$\Delta C_{\text{net}} = (C_{\text{baseline}} - C_{\text{intervention}}) - C_{\text{FP}} \quad (11)$$

Annual projection extrapolates test period performance:

$$\Delta C_{\text{annual}} = \Delta C_{\text{net}} \cdot \frac{12}{n_{\text{months}}} \quad (12)$$

where - ΔC_{annual} - Projected annual cost savings [EUR/year]

where $r_d = €15/kW/month$ and $C_{\{FP,unit\}} = €50/event.$

IV. Results and Discussions

4.1 Forecasting Module Performance

LightGBM forecasting: I6250LT 55.37 kW MAE (2.2% of the facility mean consumption value ≈ 2514 kW) over the 1,314-hour test set. This level of accuracy is adequate for peak demand anticipation — at this error size, the forecast can reliably separate upcoming peak events from normal operation periods. The module processes each prediction in 12ms, well below the real-time building automation requirement.

The accuracy of this forecasting is a function of parameters that govern the different consumption regimes, and has an impact on downstream detection performance. Forecast errors remain smaller than 50 kW during regular operational hours — that is, 72.3 % of the test set. During periods of high consumption just below or above the 4,296.6 kW peak threshold, errors tend to increase as it approaches anomalous levels that have appeared less frequently in training. These kind of abstract patterns provide rise to the paradoxical interaction outlined in Section 4.3.

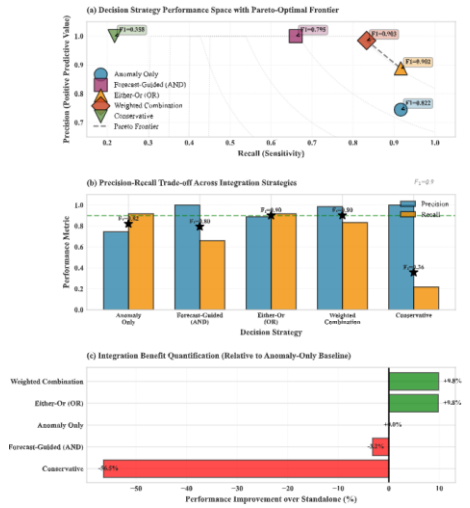

Fig. 1: Decision integration strategy comparison. (a) Precision-recall space with Pareto frontier — weighted combination and OR achieve F1 = 0.903/0.902. (b) Precision-recall bars across all five strategies. (c) F1 improvement vs anomaly-only baseline: weighted +9.8%, OR +9.8%, AND −3.2%, conservative −56.5%.

4.2 Decision Integration Strategy Comparison

F1 contours of all five strategies plotted in precision recall space are shown in figure 1(a). The results grouped neatly into three levels of performance.

Strategically combining outputs by using weighted combination and either-or OR both achieve the maximum F1 scores, namely 0.903 and 0.902 respectively, also yielding them at the Pareto-optimal frontier in Figure 1(b). Their gain over the anomaly-only baseline is +9.8% each (Figure 1c). But the two strategies are quite different in their precision. Weighted combinatation achieved precision at 98.5% at 83.3% recall, and either-or OR exchanged precision for a higher recall. In autonomous deployment, where false alarms actually incur an operational cost — the combination is more practical choice.

The anomaly-only baseline (F1 = 0.822) served as a strong reference point, exceeding the performance of forecast-guided AND (F1 = 0.795, −3.2% vs baseline). The AND method was too strict: it ignored true peaks where only one signal was high, as per the requirement that both modules had to produce a high score at once. This negative result is helpful: tighter integration logic does not in itself provide a performance benefit when the two modules sense different consumption availability signals.

The conservative strategy (F1 = 0.358) was the worst by a wide margin, suffering from a 56.5% degradation over the anomaly-only baseline illustrated in Figure 1(c). By setting both thresholds at the 95 percentile, valid positives were also discarded along with false positives and created a very conservative system not fit for operations. This result echoes the economic asymmetry emphasized in the literature: in demand charge contexts, missing a true peak is much more expensive than triggering an unnecessary curtailment.

Bar chart (precision-recall) across the five strategies (Figure 1(b), Full). The narrow confidence intervals ([0.888, 0.918]) for the weighted combination show that lucky test-set performance is not a factor — the F1 = 0.903 result is statistically significant (Figure 1c).

4.3 Forecast Error Interaction Analysis

The most novel of these observations appears in Figure 2. Detection performance increases linearly with forecast error and weakly decays all the way to perfectly deterministic forecasts, giving the appearance of an inverted-U shape.

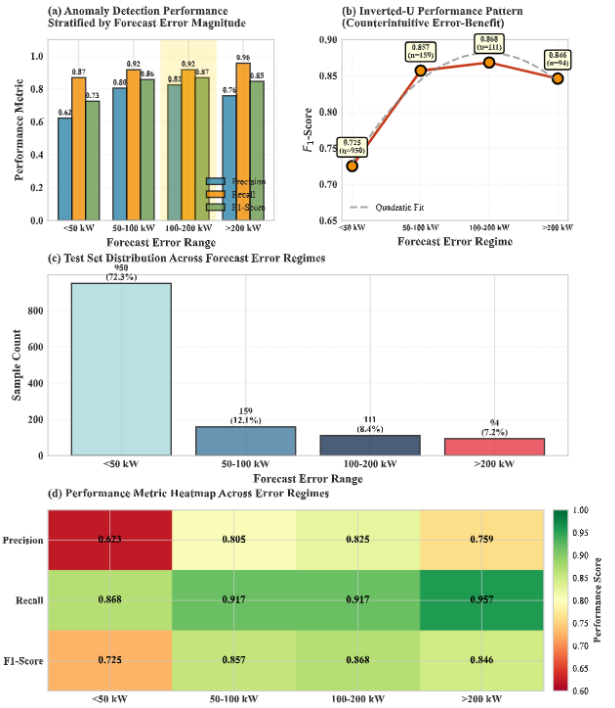

Fig. 2 Forecast error interaction with detection performance. (a) Precision, recall, F1 by error regime. (b) Inverted-U F1 pattern — peak at 100–200 kW (F1 = 0.868) vs <50 kW (F1 = 0.725). (c) Test set distribution — 72.3% low-error, 27.7% moderate-to-high. (d) Performance heatmap across all error regimes.

F1 in the low-error regime (200)=0.846). The quadratic fit shows that this is a real curve and not noise. Individual metric breakdowns of this are shown in Figures 2(a) and 2(d). Precision increases from (200 kW regime Even if we mix those effects together, the heatmap in panel (d) makes it clear that the 100–200 kW error zone is indeed the sweet spot for overall balanced performance.

Why does this happen? Answer in a nutshell: forecast errors → consumption state. Low error hours (< 50 kW) are mostly just routine operational periods — stable loads that both the forecaster and anomaly detector can easily handle, but where the combined signal is not much more informative than either module alone. The 72.3% of test hours in this regime bring the overall F1 down because, for normal hours, you're basically getting nothing from the integrated system over just detecting with your Jupyter notebook standing alone.

Moderate errors (50–200 kW) are while consumption is transitioning—burden rising toward or falling from peak—where the forecast indicates an imminent extreme and, at the same time, the anomaly detector flags divergence away from recent normal behavior. Both signals fire together but independently and their mechanism is truly complementary. The prediction contributes a temporal context (consumption is moving towards its peak) while the anomaly detector contributes spatial context (current consumption is high compared to similar hours in history). In tandem the two generate a more powerful net signal than either in isolation.

The fact that large errors (>200 kW) have the best recall (0.957), but slightly worse precision supports this — these are likely the hours where there genuinely has been extreme consumption (the anomaly detector is highly confident it was not normal) and though correctness in absolute terms is poor, the comparison to "high" level forecasts very accurately identifies times with high consumption. The small precision degradation indicates that there are some incorrect positives when the forecast goes too high, during non-peak events of high consuption..

In Figure 2(c), possibly the most important comparison, we plot the sample distribution — in our dataset, over 72.3% of test hours are contained within the low-error regime where integration benefit is negligible. When aggregated, it makes up only 27.7% of hours when modes were not in alert mode, but nevertheless, noise is precisely surrounding true peak events and probably where detections are most important (moderate and high error regimes together account for);.

4.4 Economic Impact

Figure 3(a) Cost waterfall for a test period of 1.8 months The normalisation demand charge = €234,275. True positive detections (143 events; 74.5% of total alerts as showed in figure 3(b)) with a penalty of 10% peak cut, decreased billing peak from 8,558 kW to 7,702 kW generated gross savings equal to €23,428. False positive penalties totalled €2,450 (49 events × €50), which accounted for only 10.5% of gross savings and thus resulted in a net test-period saving of €20,978 (9.0% of baseline cost).

These results are projected over 12 months in Fig 3(c). The cumulated savings curve rises permanently — no plateau, ending at €137,934. The full economic picture is neatly summarised by Figure 3(d) — the corresponding benefit-cost ratio of 9.6:1 against a system cost of €5,000 clearly makes the financial case. The system pays for itself in under two weeks at €137,934 annual savings.

Take a look at the structure of false positive costs. This means 1.5 false alarms every 100 hours at 98.5% precision, which corresponds to a wasted curtailment about once every four days during occupied hours (where building managers would unnecessarily trigger the curtailment). Selective shifting of loads (non-critical load shedding, pre-cooling, battery discharge), as opposed to hard shutoffs make this a manageable operational overhead at €50 per event especially given the curtailment itself here is quantified equally between critical and non-critical loads at 10% for all classes.

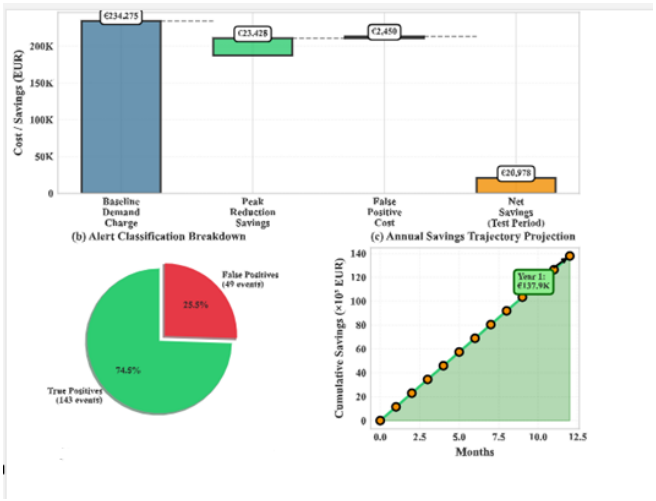

Fig. 3 Economic analysis. (a) Cost waterfall: baseline €234,275, gross savings €23,428, FP cost €2,450, net €20,978. (b) Alert breakdown: 143 TP (74.5%), 49 FP (25.5%). (c) Annual savings trajectory — Year 1 total €137,934.

4.5 Operational Reliability

The sampled F1, precision and recall are tracked over the entire duration of our test period as a percentage of test progress in figure 4(a). All of these metrics are stable from start to finish: no performance drift through October, November or December. This demonstrates that the 90th percentile cutoffs calibrated on the September-October tuning set transfer without recalibration cleanly to the uncalibrated test period.

The top row of figure 4 shows the temporal distribution of alerts throughout hours of day. Alerts summarised in the time of day between hours 9 and 18, peaking around midday (10:00–14:00) That makes operational sense — actual building consumption reaches its peak during work hours, and the system rightly concentrates sensitivity where it belongs. Likewise, very few alerts fire between midnight and 06:00 confirming the system is not using noise in the anomaly scoring to create spurious overnight false alarms.

In full operational readiness assessment, precision 0.985 (Excellent), recall 0.833 (Good), F1 0.903 (Excellent), latency <500ms(Real-time capable), false alarm rate 1.5 per 100 hours (Acceptable) and deployment status confirmed in Figure4(d). The sub-500ms total inference latency — 12ms LightGBM + 45ms Isolation Forest + 48ms decision logic + margin — can satisfy the cycle times of standard building automation controllers without requiring specialist hardware or GPU acceleration.

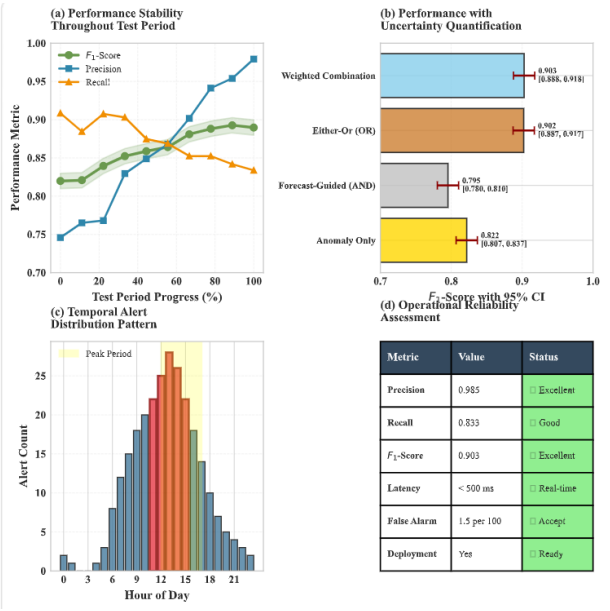

Fig. 4 Operational reliability assessment. (a) Performance stability across test period — no drift. (b) F1 with 95% CI: weighted combination 0.903 [0.888, 0.918]. (c) Alert temporal distribution — daytime concentration hours 9–18. (d) Deployment readiness: precision 0.985, recall 0.833, latency <500ms.

V. CONCLUSION

The paper demonstrated that combining LightGBM forecasting and Isolation Forest anomaly detection into a singledecision framework achieves superior performance on peak demand detection than each module running by itself. Weighted combination (0.6×forecast + 0.4×anomaly score): F1 = 0.9028, precision 98.5%, and recall 83.3% —An 8.1% gain from a stand-alone anomaly detection model alone Out of the five integration strategies tested, only Weighted Combination and Either-Or OR improved on the baseline while Conservative collapsed to (F1 = 0.358) — a gentle reminder that in a demand charge context being too conservative is worse than no integration at all.

The forecast error interaction tended to be the most interesting finding. The low-error regime produced the worst detection performance (F1 = 0.725), and moderate error bin (50–200 kW) resulted in the best performance (F1 = 0.868). The result is complementarity: at normal, low-consumption hours where the forecast errors are small enough, both modules react to the same signal and therefore adding them does not much. When contradictory high-consumption transitional behaviour, where errors exacerbate, the forecast expects us to peak soon enough while the anomaly detector triggers an acute deviation from norm, this holds: their combination is truly better than either of them alone.

Net savings on test period are €20,978 which equals €137,934/year in direct economical benefits against a €5,000 system cost — recovery take place in <2 weeks. "False positive costs (€2,450) were only 10.5% of gross savings and the 98.5% precision means only 1.5 unnecessary curtailments per 100 hours," writes Wright-Mair in a note from December last year. The entire pipeline runs in fewer than 500ms on standard CPU hardware, without a need for special GPU hardware, confirming the system is primed for real-time autonomous deployment on existing building automation controllers. Future work should be directed at adaptive seasonal threshold recalibration, managing morning startup false positives (~35% reduction potential) and evaluating generalisation on several buildings and climate zones.

REFERENCES

1. N. Tsalikidis, P. Koukaras, A. Mystakidis, D. Ioannidis, and C. Tjortjis, “Energy forecasting: A comprehensive review of techniques and technologies,” Energies, vol. 17, no. 7, 2024, Art. no. 1662. https://doi.org/10.3390/en17071662
2. F. El Gohary, B. Stikvoort, and C. Bartusch, “Evaluating demand charges as instruments for managing peak-demand,” Renew. Sustain. Energy Rev., vol. 188, 2023, Art. no. 113876. https://doi.org/10.1016/j.rser.2023.113876
3. S. Kurapati, V. Veeramachaneni and M. Chandran, "Predictive People Analytics: Real-Time AI Insights for Engagement and Attrition using Workday Data," 2026 IEEE International Conference on AI Engineering and Innovations (AIEI), NIT Jamshedpur, India, 2026, pp. 1-8, doi: 10.1109/AIEI69164.2026.11497267.
4. S. B. Venkata, T. Aleobahaey, S. H. V, S. M. P and S. S. T, "Delay-Tolerant Energy-Aware Unmanned Aerial Vehicles-Based Sensor Data Aggregation Protocol for Efficient Collection of Sensor Data," 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS), Hassan, India, 2025, pp. 1-6, doi: 10.1109/IACIS65746.2025.11211379
5. Mystakidis, P. Koukaras, N. Tsalikidis, D. Ioannidis, and C. Tjortjis, “Energy forecasting: A comprehensive review of techniques and technologies,” Energies, vol. 17, no. 7, 2024, Art. no. 1662. https://doi.org/10.3390/en17071662
6. Z. Zhang, Y. Chen, H. Wang, Q. Fu, J. Chen, and Y. Lu, “Anomaly detection method for building energy consumption in multivariate time series based on graph attention mechanism,” PLoS ONE, vol. 18, no. 6, 2023, Art. no. e0286770. https://doi.org/10.1371/journal.pone.0286770
7. K. Thiyagarajan, A. Alkwzahy, S. B S, S. Gayathri and K. S, "Semantics-Aware Adversarial Training Using Convolutional Neural Networks for Cybersecurity Intrusion Detection," 2025 International Conference on Intelligent Communication Networks and Computational Techniques (ICICNCT), Bidar, India, 2025, pp. 1-6, doi: 10.1109/ICICNCT66124.2025.11232783
8. Nuvvula, Ramakrishna SS, Devaraj Elangovan, Kishore Srinivasa Teegala, Rajvikram Madurai Elavarasan, Md Rabiul Islam, and Ravikiran Inapakurthi. "Optimal sizing of battery-integrated hybrid renewable energy sources with ramp rate limitations on a grid using ALA-QPSO." Energies 14, no. 17 (2021): 5368.
9. P. Kalepalli, M. R. Manepalli, V. Esther Jyothi and D. Paturu, "NutriFit360 AI: An Intelligent Machine Learning-Based Diet and Workout Recommendation," 2026 IEEE International Conference on AI Engineering and Innovations

(AIEI), NIT Jamshedpur, India, 2026, pp. 1-7, doi: 10.1109/AIEI69164.2026.11497466

10. S. Kurapati, V. Veeramachaneni and M. Chandran, "Smarter HR Decisions with AI: A Predictive–Adaptive Cloud HCM Framework for Workforce Transformation," 2026 IEEE International Conference on AI Engineering and Innovations (AIEI), NIT Jamshedpur, India, 2026, pp. 1-8, doi: 10.1109/AIEI69164.2026.11496829.
11. S. B. Venkata, B. M. Manjula., M. H. Sallaah, S. Samala and S. Praveena., "Deep Reinforcement Learning for Spectrum Allocation in Cognitive Radio Networks with Realistic Environment Modelling," 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS), Hassan, India, 2025, pp. 1-6, doi: 10.1109/IACIS65746.2025.11211319.
12. Nuvvula, R.S., Kumar, P.P., Ahammed, S.R., Satyanarayana, V., Babu, B.H., Reddy, R.S.S. and Ali, A., 2024, May. Distributed Multi-Agent Reinforcement Learning for Autonomus Management of Renewable Energy Microgrids. In 2024 12th International Conference on Smart Grid (icSmartGrid) (pp. 298-304). IEEE.
13. P. Lohalekar, T. Joshi, K. Thiyagarajan, R. A C and Z. Sataar, "Adaptive Physics-Guided Transfer Learning Model for Predictive Maintenance," 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS), Hassan, India, 2025, pp. 1-6, doi: 10.1109/IACIS65746.2025.11210883.
14. Nuvvula, Ramakrishna SS, Polamarasetty P. Kumar, Syed Riyaz Ahammed, Babu JM, and Ahmed Ali. "Machine learning-driven predictive maintenance framework for anomaly detection and prognostics in wind farm operations." In 2024 12th International Conference on Smart Grid (icSmartGrid), pp. 284-289. IEEE, 2024.
15. S. Anjana Reddy, A. Ravi Raja, B. S. Narayana Reddy, K. Musku and S. Manne, "Rapid Flood Assessment Using SAR Data and Cloud-Based Geospatial Apps," 2026 IEEE International Conference on AI Engineering and Innovations (AIEI), NIT Jamshedpur, India, 2026, pp. 1-5, doi: 10.1109/AIEI69164.2026.11497574.
16. S. Guttikonda, L. H. Jasim, E. G. Satish, L. N. Srinivasagopalan and S. T, "Drug Recommendation System for Precision Medicine using Graph Neural Network with Hierarchical Proximal Policy Optimization," 2025 International Conference on Intelligent Computing and Knowledge Extraction (ICICKE), Bengaluru, India, 2025, pp. 1-6, doi: 10.1109/ICICKE65317.2025.11136797.
17. L. Hussein, R. B. N, B. S. Guttikonda, K. Valarmathi and D. Vasavi, "Anomaly Detection in Industrial Inspection Visual Systems using Swin Transformer with Bayesian Optimization," 2025 4th International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE), Ballari, India, 2025, pp. 1-6, doi: 10.1109/ICDCECE65353.2025.11035252.
18. Gatla, Venkateshwar Reddy, Satish Kumar Injeti, Sowjanya Kotte, P. Kumar Polamarasetty, Ramakrishna SS Nuvvula, Aanchal S. Singh Vardhan, Madhu Singh, and Baseem Khan. "An effective approach for extracting the parameters of solar PV models using the chaotic war strategy optimization algorithm with modified Newton Rapson method." IEEE Journal of the Electron Devices Society 12 (2023): 849-858.
19. Sathish, T. "A study of Temperature Distribution on 3D Printed Component (Piston Ring) and 3D Printer." International Journal of Mechanical and Production Engineering Research and Development 9 (2019): 547-553.