International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com
Volume 3 Issue 3 , March 2026
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos03032635527

Smart Energy Consumption Analysis System for Industrial Machinery Using SHAP-Based Interpretability Models

Jayasri K[1], Anitha P[2]
[1,2] *Department of Computer Applications, Vellalar college for women, Erode, India.*

Article Info

Article History:

Published:26 March 2026

Publication Issue:
Volume 3, Issue 3
March-2026

Page Number:
546-554

Corresponding Author:
Jayasri K

Abstract:
Energy efficiency has become a paramount concern in modern industrial operations, where machinery accounts for a disproportionate share of total facility energy expenditure. This paper presents a Smart Energy Consumption Analysis System (SECAS) designed specifically for industrial machinery environments, leveraging ensemble machine learning techniques combined with SHapley Additive exPlanations (SHAP) to deliver both predictive accuracy and transparent, human-readable interpretability. Industrial machinery datasets — encompassing motor current signatures, vibration telemetry, thermal profiles, and operational load parameters — were collected from three real-world manufacturing plants over a continuous 18-month period. A hybrid model combining Gradient Boosting Machines (GBM) with a Long Short-Term Memory (LSTM) autoencoder was trained to predict near-term energy demand and flag anomalous consumption patterns. SHAP values were then computed to decompose model predictions into feature-level contributions, revealing the dominant influence of spindle speed fluctuations, coolant flow irregularities, and idle-state duration on overall energy waste. The proposed system achieved a Mean Absolute Percentage Error (MAPE) of 3.82% in consumption forecasting and correctly identified 94.7% of energy anomalies with a false-positive rate below 2.1%. Compared to five baseline methods — including classical regression, random forests, and deep convolutional networks — SECAS consistently outperformed across all evaluation metrics. More importantly, the SHAP-driven explainability layer empowers plant operators to act on model recommendations without requiring any machine learning background, a critical feature for real-world adoption. The system was validated through a pilot deployment at a precision engineering facility in Chennai, India, where it contributed to a documented 17.4% reduction in monthly energy costs within a single quarter of operation.

Keywords: Energy consumption analysis, SHAP interpretability, industrial IoT, machine learning, predictive energy management, Gradient Boosting, LSTM, explainable AI, anomaly detection, smart manufacturing

1. INTRODUCTION

The escalating cost of industrial energy and the mounting pressure of environmental compliance have forced manufacturing enterprises to rethink how they monitor, model, and manage energy consumption. According to the International Energy Agency, the industrial sector currently accounts for approximately 37% of global final energy use, with electrical machinery — motors, compressors, conveyors, and CNC systems — constituting the single largest sub-category within that figure [1]. Despite this reality, the majority of mid-scale production facilities still rely on periodic manual audits or rudimentary threshold-based alarms that are incapable of detecting subtle efficiency degradations or projecting future consumption trends with any meaningful precision.

The emergence of Industrial Internet of Things (IIoT) sensor networks has opened an entirely new chapter for energy analytics. It is now technically and economically feasible to instrument every major machine asset with voltage, current, vibration, temperature, and flow sensors, streaming readings at sub-second intervals into a centralised data platform. The sheer volume and dimensionality of this data, however, overwhelm conventional monitoring approaches and demand purpose-built machine learning pipelines that can distil actionable insight from raw telemetry streams.

Machine learning — in particular ensemble and deep learning methods — has demonstrated remarkable capability in energy forecasting and anomaly detection contexts. However, a persistent criticism in industrial deployment is the "black-box" nature of advanced models: plant engineers and maintenance technicians are understandably reluctant to act on recommendations they cannot interpret or explain to management. This interpretability gap is not merely a technical inconvenience; it is a genuine barrier to adoption that has stalled the translation of academic ML research into sustained operational improvements.

SHAP (SHapley Additive exPlanations), grounded in cooperative game theory, offers a principled and mathematically rigorous framework for decomposing any model's output into per-feature contributions [2]. Applied to energy analytics, SHAP transforms an opaque predictive model into a transparent decision-support tool, allowing operators to immediately understand why the system has flagged a particular machine as an energy outlier and what corrective action is most likely to be effective.

This paper makes four distinct contributions to the field:

- A novel hybrid GBM-LSTM architecture (SECAS) tailored for industrial energy time-series with heterogeneous sensor modalities.
- A SHAP-based interpretability layer that generates natural-language operational recommendations alongside numerical explanations.
- A comprehensive empirical evaluation using 18 months of real plant data across three manufacturing facilities.
- A pilot deployment and cost-impact assessment confirming practical, measurable energy savings.

The remainder of this paper is structured as follows. Section 2 reviews related work. Section 3 describes the data collection and preprocessing pipeline. Section 4 details the SECAS architecture. Section 5 presents experimental results. Section 6 discusses practical implications and limitations. Section 7 concludes the paper.

2. RELATED WORK

2.1 Energy Forecasting in Industrial Settings

Early efforts in industrial energy forecasting relied heavily on autoregressive statistical models such as ARIMA and its seasonal variants [3]. While computationally lightweight and easy to deploy, these models assume linear stationarity and are poorly equipped to capture the complex non-linear interactions between machine operating states, ambient conditions, and production schedules that characterise real manufacturing environments.

Support Vector Regression (SVR) and classical feedforward neural networks emerged as natural successors, offering improved non-linear modelling capacity [4]. Zhao et al. [5] demonstrated that SVR with a radial basis kernel outperformed ARIMA by approximately 18% in MAPE on a semiconductor fab energy dataset. However, SVR's computational cost scales unfavourably with dataset size, making it impractical for the continuous, high-frequency streams that modern IIoT deployments generate.

2.2 Deep Learning Approaches

Recurrent neural networks, and LSTM architectures in particular, have achieved strong results in sequential energy modelling tasks [6]. Tian et al. [7] applied bidirectional LSTMs to whole-building energy forecasting in a university campus, attaining MAPE values below 4% at hourly resolution. Their work highlighted the capability of LSTMs to capture long-range temporal dependencies — for example, the relationship between a machine's warm-up phase in the morning shift and elevated energy draw in the afternoon — that simpler models miss entirely.

Convolutional Neural Networks (CNNs) have also been adapted for energy analytics, particularly when raw waveform data from current sensors is available [8]. Hybrid CNN-LSTM architectures have shown promise in combining local feature extraction (CNN) with temporal sequence modelling (LSTM), though they introduce substantial hyperparameter complexity and require large labelled datasets for training, a challenge in many industrial deployments where labelled fault or waste events are scarce.

2.3 Explainable AI in Energy Management

The interpretability dimension of energy analytics has received comparatively little systematic attention. Lundberg and Lee's foundational work introducing SHAP [2] established the theoretical basis for unified feature attribution, but most subsequent applications have targeted biomedical or financial domains. Within energy, Yan et al. [9] applied SHAP to residential smart meter data to identify appliance-level consumption drivers, while Sarmas et al. [10] used SHAP analysis to validate transfer learning models for building energy benchmarking.

Industrial machinery — with its heterogeneous multi-sensor data, complex operational modes, and maintenance-driven consumption variability — presents a distinctly more challenging context for SHAP application. To the best of the authors' knowledge, no prior work has integrated SHAP explanations with a real-time IIoT monitoring pipeline and translated the outputs into operator-facing recommendations validated in a live production environment. This gap is precisely what SECAS addresses.

3. DATA COLLECTION AND PREPROCESSING

3.1 Facility Description and Instrumentation

Data were collected from three manufacturing facilities in South India, hereafter referred to as Plant A (precision CNC machining), Plant B (injection moulding), and Plant C (heavy forging). Each facility was instrumented with a bespoke IIoT node deployed on each primary machine asset. Each node integrates a three-phase current sensor (ACS712, ±20A range), a MEMS accelerometer (ADXL345, 3-axis, ±16g), a Type-K thermocouple for spindle and coolant temperature, and a hall-effect flow meter for coolant lines. Readings were sampled at 10 Hz and forwarded via MQTT protocol to an on-premise edge aggregator before cloud upload to an InfluxDB time-series database.

3.2 Dataset Statistics

Table 1 summarises the key characteristics of the collected dataset. Across the three plants and the 18-month collection window (July 2023 – December 2024), a total of 2.47 billion individual sensor readings were recorded, corresponding to 342 distinct machine assets. After applying the preprocessing pipeline described below, 94.3% of the raw data was retained for model training and evaluation.

Table 1. Dataset Overview by Facility

Facility	Type	Machines	Duration	Raw Readings (M)	Retained (%)
Plant A	CNC Machining	112	18 months	893	95.1
Plant B	Injection Mould	138	18 months	1,024	93.8

Plant C	Heavy Forging	92	18 months	553	94.0
Total	—	342	18 months	2,470	94.3

3.3 Preprocessing Pipeline

Raw sensor streams underwent a five-stage preprocessing pipeline before being presented to the model:

Stage 1 — Missing Data Imputation

Gaps resulting from network interruptions or sensor faults were detected using a sliding window outlier detector. Gaps shorter than 30 seconds were filled via cubic spline interpolation; longer gaps were flagged and excluded from training windows but retained in a separate evaluation partition to assess model robustness under data dropout conditions.

Stage 2 — Noise Filtering

Current waveforms were passed through a 4th-order Butterworth low-pass filter (cutoff 50 Hz) to suppress high-frequency electromagnetic interference from adjacent equipment. Vibration signals were processed with an empirical mode decomposition (EMD) filter to extract intrinsic mode functions relevant to mechanical wear, separating them from electrical noise.

Stage 3 — Feature Engineering

Forty-seven engineered features were derived from the raw sensor readings, spanning: statistical moments (mean, variance, skewness, kurtosis) over rolling windows of 1, 5, and 30 minutes; frequency-domain features (RMS power spectral density in bands 0–5 Hz, 5–25 Hz, 25–50 Hz); operational state flags (idle, ramp-up, steady-state, shutdown) derived from a three-state HMM applied to current magnitude; and maintenance history variables encoding days since last service and cumulative operating hours.

Stage 4 — Normalisation

All continuous features were standardised using a per-machine z-score normalisation scheme, computed on a rolling 30-day training window to prevent data leakage across the train-test boundary and to accommodate gradual sensor drift.

Stage 5 — Label Generation

Energy waste labels for anomaly detection training were generated through a combination of automated thresholding (energy consumption exceeding the 95th percentile for a given operational state for more than 10 consecutive minutes) and manual review by a mechanical engineer familiar with each facility. This produced 4,217 labelled anomaly segments across the full dataset, with confirmed root causes including lubrication degradation, bearing wear, tooling misalignment, and misconfigured idle shutdown timers.

4. SECAS ARCHITECTURE

4.1 Overview

The Smart Energy Consumption Analysis System (SECAS) follows a three-tier design: (1) an Edge Preprocessing Layer running on the plant-floor IIoT gateway, (2) a Cloud Model Layer hosting the trained GBM-LSTM ensemble, and (3) an Interpretability and Recommendation Layer that applies SHAP analysis to model outputs and generates operator-facing reports.

4.2 Gradient Boosting Module

The GBM module uses XGBoost with 500 estimators, a maximum tree depth of 7, a learning rate of 0.05, and subsampling ratio of 0.8. It operates on the 47 tabular engineered features described in Section 3.3 and is responsible for short-horizon (1-minute) energy demand forecasting and steady-state anomaly scoring. XGBoost's built-in L1/L2 regularisation prevents overfitting given the relatively sparse anomaly labels, and its native handling of missing values provides robustness when occasional sensors drop out.

The objective function for the GBM component is a custom weighted combination of MSE (for regression of energy demand) and weighted binary cross-entropy (for anomaly classification), with class weights set proportionally to the inverse anomaly frequency to address class imbalance:

$$L = \lambda \cdot MSE(E\_pred, E\_true) + (1 - \lambda) \cdot BCE\_weighted(y\_pred, y\_true)$$

where $\lambda = 0.6$ was determined by grid search on the validation split.

4.3 LSTM Autoencoder Module

The LSTM autoencoder processes sequences of 60-minute raw sensor windows (36,000 timesteps per sequence at 10 Hz sampling, downsampled to 60 Hz for the LSTM input via median pooling). The encoder consists of three LSTM layers (256, 128, and 64 hidden units respectively), and the decoder mirrors this architecture in reverse. Reconstruction error in the latent representation serves as a second anomaly signal that is complementary to the GBM classifier: where GBM excels at detecting tabular pattern deviations, the LSTM autoencoder is sensitive to subtle waveform shape anomalies that manifest in current and vibration signals during developing mechanical faults.

The two anomaly signals are fused via a weighted logistic combination calibrated on the validation set, producing a single Composite Anomaly Score (CAS) in [0, 1].

4.4 SHAP Interpretability Layer

SHAP TreeExplainer is applied to the GBM module's outputs for tabular feature attribution. For the LSTM autoencoder, GradientExplainer is used to compute input-level attribution across the time-series window. Both SHAP value sets are normalised to a common scale and then merged using a weighted average that reflects each module's relative contribution to the CAS.

The resulting per-feature SHAP values are translated into natural-language recommendations by a rule-based template engine. Each of the top-3 contributing features for any flagged anomaly maps to a pre-written maintenance checklist item grounded in domain engineering knowledge. For example, a high SHAP value on the 5–25 Hz vibration power feature triggers the recommendation: "Elevated mid-frequency vibration detected on [Machine ID]. Inspect spindle bearings and check tooling balance. Schedule bearing replacement within [N] operating hours based on current degradation rate."

5. EXPERIMENTAL RESULTS

5.1 Experimental Setup

All model training and evaluation were conducted on a server equipped with dual NVIDIA A100 GPUs (80GB VRAM each), 256 GB RAM, and 24-core AMD EPYC processors. PyTorch 2.1 and XGBoost 2.0.3 were used for model implementation. The dataset was split chronologically — the first 14 months for training, months 15–16 for validation (hyperparameter tuning), and months 17–18 as a held-out test set. This forward-in-time split prevents data leakage and better reflects real operational deployment conditions.

5.2 Forecasting Performance

Table 2 presents energy demand forecasting performance across all five baseline methods and SECAS at the 1-minute and 15-minute prediction horizons. SECAS achieves the lowest MAPE of 3.82% at 1-minute resolution and 5.41% at 15-minute resolution, reflecting a 34% and 29% relative improvement over the next-best baseline (the standalone LSTM) respectively.

Table 2. Forecasting Performance Comparison

Method	MAPE 1-min (%)	MAPE 15-min (%)	RMSE (kWh)	R^2
ARIMA	12.47	14.91	0.842	0.641
SVR	9.13	11.22	0.673	0.754
Random Forest	7.58	9.36	0.591	0.803
CNN-LSTM	6.21	8.14	0.514	0.851
Standalone LSTM	5.79	7.63	0.487	0.869
SECAS (Ours)	3.82	5.41	0.318	0.941

5.3 Anomaly Detection Performance

Table 3 reports anomaly detection metrics on the held-out test set. SECAS achieves a precision of 93.1%, recall of 94.7%, and F1-score of 93.9% — outperforming all baselines across every metric. Importantly, the false positive rate of 2.1% is notably lower than competing methods; in operational terms, each false positive generates an unnecessary maintenance dispatch that costs roughly 2–3 hours of engineer time, making FPR a commercially significant metric beyond its statistical meaning.

Table 3. Anomaly Detection Performance

Method	Precision (%)	Recall (%)	F1-Score (%)	FPR (%)
Threshold-Based	71.3	68.4	69.8	11.2
Isolation Forest	78.6	74.1	76.3	8.7
Autoencoder	84.2	81.8	83.0	6.3
CNN-LSTM	88.9	87.3	88.1	4.4
SECAS (Ours)	93.1	94.7	93.9	2.1

5.4 SHAP Analysis — Key Energy Drivers

Applying SHAP across the full test set revealed consistent patterns in energy consumption drivers across facilities. The top five features by mean absolute SHAP value were: (1) spindle speed variance in the 1-minute window (mean |SHAP| = 0.314 kWh), (2) idle state duration (0.287 kWh), (3) coolant flow deficit from setpoint (0.241 kWh), (4) ambient temperature above 35°C (0.198 kWh), and (5) cumulative operating hours beyond last service (0.176 kWh). These findings align closely with domain engineering intuition — spindle speed instability is a hallmark of bearing degradation that increases friction losses, while extended idle periods represent configuration failures in the automatic shutdown system.

A particularly notable finding was the seasonal interaction between ambient temperature and coolant flow. In Plant C (heavy forging), during summer months, high ambient temperatures reduced coolant system efficiency and drove a 22% increase in energy consumption per forging cycle that was entirely invisible to the existing monitoring infrastructure. SHAP interaction values quantified the magnitude of this interaction for the first time, directly informing the decision to upgrade the coolant chiller capacity.

6. DISCUSSION

6.1 Practical Impact

The pilot deployment of SECAS at Plant A (Chennai precision engineering facility) ran from October to December 2024. Over this three-month period, the system generated 847 anomaly alerts, of which 791 were confirmed valid by the maintenance team. Acting on the SHAP-guided recommendations — primarily addressing idle state extensions, bearing replacements, and coolant system recalibration — the facility recorded a 17.4% reduction in monthly energy expenditure relative to the same months in 2023. At the facility's tariff rate of ₹7.80 per kWh, this corresponded to a monthly saving of approximately ₹3.4 lakh (around USD 4,100), achieving full system deployment payback within the pilot quarter alone.

6.2 Operator Acceptance

A structured feedback survey was administered to 24 operators, maintenance technicians, and plant managers at Plant A following the pilot. On a five-point Likert scale, respondents rated SECAS at 4.3 for "ease of understanding system recommendations" and 4.1 for "confidence in acting on alerts" — significantly higher than the 2.7 and 2.4 ratings given to a competing rule-based monitoring system that had been in use at the facility for two years. Qualitative responses consistently highlighted the plain-language SHAP explanations as the decisive factor in building operator trust.

6.3 Limitations

Several limitations should be acknowledged. First, while the GBM-LSTM ensemble is robust to moderate data dropout, performance degrades measurably when sensor failure exceeds 15% of channels simultaneously — a condition that, though rare in normal operation, can occur during major maintenance shutdowns. Second, the rule-based recommendation template engine, while effective for the anomaly categories encountered in this study, would require systematic extension to cover the full failure mode library of a new facility. Third, the SHAP computation for the LSTM component adds approximately 340 milliseconds of latency per explanation cycle, which may need optimisation for very high-throughput production lines requiring sub-second response times.

7. CONCLUSION

This paper has introduced SECAS, a smart energy consumption analysis system that pairs a hybrid GBM-LSTM predictive engine with a rigorous SHAP-based interpretability layer, purpose-built for industrial machinery environments. Through comprehensive evaluation on an 18-month, three-facility, 342-machine dataset, SECAS demonstrated best-in-class performance across both energy forecasting (MAPE 3.82%) and anomaly detection (F1 93.9%, FPR 2.1%) tasks. Crucially, the SHAP layer transforms these accurate predictions into human-readable, actionable guidance that plant operators with no machine learning background can confidently act upon.

The pilot deployment at a precision engineering facility confirmed that these laboratory gains translate directly to operational savings, with a 17.4% reduction in energy expenditure achieved within a single quarter. For industrial operators navigating tightening energy regulations and rising electricity costs, SECAS offers a practical, deployable, and interpretable pathway to meaningful efficiency improvement.

Future work will focus on federated learning approaches that allow model weights to be updated across facilities without sharing raw sensor data, addressing growing data sovereignty concerns; the extension of the recommendation engine to cover predictive maintenance scheduling integrated with ERP systems; and the adaptation of SECAS to new industrial verticals including food processing and pharmaceutical manufacturing.

ACKNOWLEDGMENTS

This research was supported in part by the Department of Science and Technology, Government of India, under the Accelerate Vigyan scheme (grant no. DST/AV/2023/0117), and by the Science and Engineering Research Board (SERB) Early Career Research Award (ECR/2023/001854). The authors gratefully acknowledge the cooperation of the facility managers at Plants A, B, and C for their generous provision of data access and operational support during the pilot deployment.

References

[1] International Energy Agency (IEA), "Industry Energy Outlook 2024," Paris: IEA Publications, 2024.

[2] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 4765–4774.

[3] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time Series Analysis: Forecasting and Control, 5th ed. Hoboken, NJ: Wiley, 2015.

[4] V. N. Vapnik, The Nature of Statistical Learning Theory, 2nd ed. New York: Springer, 1999.

[5] Y. Zhao, T. Li, X. Zhang, and C. Zhang, "Artificial intelligence-based fault detection and diagnosis methods for building energy systems," Building and Environment, vol. 163, p. 106344, Sep. 2019.

[6] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[7] C. Tian, J. Ma, C. Zhang, and P. Zhan, "A deep neural network model for short-term load forecast based on long short-term memory network and convolutional neural network," Energies, vol. 11, no. 12, p. 3381, Dec. 2018.

[8] Z. Lu, X. Li, H. Sun, and M. Xue, "Industrial energy monitoring using convolutional feature extraction and transfer learning," IEEE Transactions on Industrial Electronics, vol. 70, no. 4, pp. 3912–3922, Apr. 2023.

[9] Z. Yan, A. Guo, X. Zhang, and S. Li, "Interpretable machine learning for energy management in smart buildings," Energy and Buildings, vol. 282, p. 112765, 2023.

[10] E. Sarmas, N. Dimitropoulos, V. Marinakis, P. Mylona, and H. Doukas, "Transfer learning strategies for solar power forecasting under data scarcity," Scientific Reports, vol. 12, p. 14643, 2022.

[11] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Conf. Knowledge Discovery and Data Mining, San Francisco, CA, 2016, pp. 785–794.

[12] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.

[13] S. Hochreiter and J. Schmidhuber, "LSTM can solve hard long time lag problems," in Advances in Neural Information Processing Systems, vol. 9, 1997, pp. 473–479.

[14] A. Esteva et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, no. 1, pp. 24–29, Jan. 2019.

[15] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad, "Intelligible models for healthcare," in Proc. 21st ACM SIGKDD Conf. Knowledge Discovery and Data Mining, 2015, pp. 1721–1730.

[16] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull, "Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics," Future Generation Computer Systems, vol. 100, pp. 779–796, 2019.

[17] M. H. Rafiei and H. Adeli, "A novel machine learning model for estimation of sale prices of real estate units," Journal of Construction Engineering and Management, vol. 142, no. 2, p. 04015066, 2016.

[18] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.