

A Review of Secure Edge-Assisted IoT Frameworks for Industrial Energy Management: Integrating Smart Metering, Demand-Side Control, and Adversarial Resilience

Bhumika Sham Kohokade¹, Swaleha Riyaz Sayyed², Sharmeen Javed Patel³, Diba Ansari⁴, and Bhimrao Dabhade⁵

¹²³⁴⁵Department of Electrical Engineering, Guru Gobind Singh College of Engineering and Research Centre, Nashik, Maharashtra, India

Abstract - Industrial facilities need energy management systems that are measurable, controllable and trustworthy at the same time, yet the literature tends to treat these concerns separately. This paper synthesises five recent studies on IoT-based energy management, covering smart-meter deployment, edge-assisted analytics, power-quality-aware demand-side control, predictive management under adversarial attack, and advanced smart-grid technologies. From them it derives a layered, secure edge-assisted framework for industrial energy management. The framework combines three-phase metering, rule-engine analytics at the edge, automated priority-based load control with power factor correction, LSTM forecasting, heartbeat-based health monitoring, and defences against false data injection and adversarial perturbation. We formalise the analytics, control and attack models, map threats to countermeasures, and consolidate the quantitative results reported by the source studies. Those results include a 41.58% reduction in apparent power through automatic power factor correction, a 23.46% reduction in peak real power through automated load shedding, edge CPU utilisation of about 5.6% for 19 metered assets, and forecasting accuracy that falls to roughly 30–40% under gradient-based attacks and recovers to above 80% with defences in most cases. We also identify open gaps in transport security, heartbeat integrity, data quality and scalability. The paper is a conceptual and comparative contribution: the framework is not validated experimentally here, and a testbed plan is provided.

Index Terms - Industrial IoT, energy management system, edge computing, demand-side management, smart meter, adversarial attacks, cyber-physical security.

I. INTRODUCTION

Industrial electricity demand is rising, and the pressure falls hardest on peak operating hours, when aggregated production processes, auxiliary systems and shift patterns overlap [5]. Internet of Things (IoT) technology answers part of this problem. Networked meters, microcontrollers and edge servers can measure consumption at fine granularity,

process it close to the source, and act on it automatically [2], [3]. The same connectivity enlarges the attack surface. An energy management system (EMS) that forecasts supply from weather data and dispatches batteries on that forecast can be pushed into supply-demand imbalance by corrupted inputs [1].

The five studies reviewed here each solve one part of this problem. Saleem et al. [3] deploy IoT smart meters and a middleware platform across four commercial buildings. Mirani et al. [2] move analytics to an edge server on a factory floor. Tweve et al. [5] show that power-factor-aware automated load control improves peak characteristics without relying on voluntary user response. Asef et al. [1] secure a predictive EMS against adversarial attacks. Younis et al. [4] survey 5G, fog computing, heartbeat monitoring, self-healing and multi-agent reinforcement learning, and call for integrated frameworks. No single study combines metering, edge analytics, active control and security.

The contributions of this paper are: (i) a comparative synthesis of the five studies and their stated limitations; (ii) a layered secure edge-assisted framework for industrial energy management with explicit analytics, control and attack models; (iii) a threat-to-countermeasure mapping; (iv) a consolidated table of the quantitative results the source studies report; and (v) a list of gaps and a validation plan. The paper is organised as follows. Section II reviews related work, Section III presents the framework, Section IV the security model, Section V the reported evidence, Section VI discusses limitations and future work, and Section VII concludes.

II. RELATED WORK

A. IoT Smart Metering and Middleware

Saleem et al. [3] build a smart meter from three microcontrollers: an MSP430F67641 that computes voltage, current, active, reactive and apparent power, power factor and frequency; a TM4C129x that stores data and manages

control; and a NodeMCU that publishes JSON over MQTT to a middleware server. Data are sent about every 30 s, and consumers read them through an API and the Entrack application. Installed at the main distribution panels of four facilities of one company, the system exposed weekday patterns, seasonal air-conditioning effects and half-hourly demand profiles. The work is monitoring-centred; control of connected devices is named as an extension, and security is not examined.

B. Edge-Assisted Industrial Monitoring

Mirani et al. [2] argue that earlier IoT energy monitors rely on prototype hardware and underuse edge computation. Their system connects 18 wired channels (three eight-channel meters) and one wireless node on a mobile welding machine to a ThingsBoard Edge instance over MQTT, at one-minute sampling. A JavaScript rule engine derives nine parameters from the cumulative kWh counter alone: per-minute, daily, 24-hour and weekly-average energy, percentage deviation from the weekly average, and four fleet-level totals. An interval-based fetch rule tolerates power or network outages. Only the wireless meter's filtered data go to the cloud, to reduce cost. The authors note that the meters publish over unencrypted MQTT on port 1883 and defer stronger security to future work.

C. Power-Quality-Aware Demand-Side Management

Tweve et al. [5] note that demand-response schemes rely on voluntary participation and usually aggregate energy metrics without separating active, reactive and apparent power. Their framework uses a PZEM-004T sensor, an ESP32 and relays to classify loads as critical or non-critical, shed non-critical loads automatically, and correct power factor with a capacitor bank. MQTT and HTTPS link the

controller to a web interface. The evaluation combines a Proteus simulation, a calibrated hardware prototype and historical load data from an operating industrial facility. Their reading is that power factor correction relieves current and inverter capacity but does little for real energy use, while load shedding does reduce real peak demand.

D. Predictive EMS and Adversarial Security

Asef et al. [1] (SIEMS) forecast PV supply one day ahead with LSTM networks fed by meteorological and microgrid data, and use the forecast to keep battery state of charge high. They attack the training data with the fast gradient sign method (FGSM), the basic iterative method (BIM) and DeepFool, and defend with adversarial training and defensive distillation. Their detection design uses two observation layers: the day-ahead prediction and online observation at intelligent electronic devices. This is the only reviewed study that quantifies the effect of adversarial data on an EMS.

E. Advanced Smart-Grid Technologies

Younisse et al. [4] review five technologies: 5G, fog computing, heartbeat monitoring, self-healing, and multi-agent reinforcement learning (MARL). They conclude that most fog-based security work targets confidentiality and integrity while neglecting availability and node compromise; that heartbeat studies seldom consider heartbeat spoofing or jamming; that self-healing work assumes ideal communication; and that no study integrates the five technologies. They also note reliance on outdated or non-grid datasets such as NSL-KDD. Table I summarises the five studies.

TABLE I LITERATURE REVIEW OF THE FIVE REVIEWED PAPERS

AUTHOR (YEAR)	Purpose	Methodology	Key Findings	Limitations	Notes
M. U. Saleem et al. (2021)	To design, implement, and validate a low-cost IoT-based Smart Energy Management System (SEMS) for real-time electricity monitoring and energy conservation.	Developed a smart meter using MSP430F67641, TM4C129x and NodeMCU. MQTT and IoT middleware were used for data transmission. The system was installed and evaluated through an industrial case study.	The system enabled real-time monitoring of electricity consumption and two-way control of connected devices. It helped identify unnecessary energy use and supported energy conservation.	The study mainly focused on monitoring. Advanced internet-based control, smart-grid configuration, and ML/DL-based prediction were suggested as future extensions.	Practical IoT-based implementation with an industrial case study.
M. Ullah et al. (2022)	To propose a conceptual Industrial Energy Management System (IEnMS) framework using IoT and Big Data analytics.	Literature review and survey of 10 large companies. Findings were used to develop a general theoretical framework for industrial energy management.	IoT and Big Data can improve real-time energy monitoring, energy efficiency, forecasting, predictive maintenance, decision-making, and productivity.	The framework is high-level and conceptual. Specific IoT devices, communication protocols, and data frames were outside the scope.	Conceptual framework; no detailed hardware implementation.
P. Asef et al. (2022)	To improve the cybersecurity of IoT-based intelligent energy management systems against adversarial attacks.	Used an IoT-enabled grid-tied PV microgrid with BESS and LSTM-based energy prediction. FGSM, BIM and DeepFool attacks were evaluated with defensive techniques.	Adversarial attacks can significantly reduce prediction performance, while adversarial training and defensive distillation can improve robustness.	Focused on selected adversarial attacks and a specific microgrid environment. More extensive real-world and distributed security evaluation is required.	Combines AI, IoT, energy management and cybersecurity.
A. A. Mirani et al. (2024)	To develop an IoT-and edge-computing-based system for monitoring and processing industrial energy data locally.	Energy meters were connected to an edge gateway, database, rule engine and cloud platform. Multiple industrial assets were monitored and additional energy parameters were calculated at the edge.	Provided comparative industrial energy monitoring at the edge with average CPU utilization up to 5.58%. Interval-based data retrieval helped reduce measurement errors from power and network issues.	Further work is required for long-duration power failures, scalability, stronger MQTT/TLS security, and ML-based anomaly detection.	Strong practical focus on edge computing and industrial IoT.

E. Tweve et al. (2026)	To manage peak energy demand in industrial solar microgrids and improve active, reactive and apparent power performance.	Used simulation and an IoT-based prototype with real-time power-factor monitoring, capacitor-bank compensation and non-critical load shedding.	The approach reduced peak demand and improved power-factor-related performance, demonstrating the potential of automated demand-side energy management.	Further validation is needed for scalability, different industrial load profiles, and closed-loop stability under changing operating conditions.	Recent work focused on automated demand-side energy management and solar microgrids.
------------------------	--	--	---	--	--

F. Communication Technologies and Positioning

Saleem et al. [3] compare the communication options available to a smart meter. Serial links are cheap and reliable but short-range; Zigbee suits home and building networks; GPRS/GSM and WiMAX cover wide areas at higher cost; power-line communication needs little infrastructure but suffers attenuation and noise. [2] adds that MQTT is lightweight, broker-based and interoperable, and that it suits a multi-application edge better than IEC 61850, which is aimed at substation applications [2]. The framework here therefore uses MQTT between meters and the edge, and leaves the choice of physical link (Ethernet, Zigbee, Wi-Fi or 5G) to site conditions.

Commercial industrial energy platforms are the other point of comparison. [2] reviews cloud-centred products and finds them mostly closed, with limited customisation, little support for local processing, and high cost for large historical data. An open edge-centred design with selective cloud forwarding addresses these limits, but it also shifts responsibility for security to the operator, which motivates Section IV.

III. PROPOSED FRAMEWORK

The framework has five layers and a cross-cutting security plane (Fig. 1). It keeps the elements that the reviewed studies validated individually and adds the integration they call for.




Fig. 1. Layered secure edge-assisted framework for industrial energy management.

A. Metering and Perception Layer

Wired three-phase meters monitor fixed assets, and wireless nodes with a gateway monitor mobile ones, as in [2]. Each meter reports voltage, current, active power P and reactive power Q, power factor and cumulative energy. Following [5], loads are classified as critical or non-critical by production continuity, safety and redundancy. A weather station supplies irradiance, temperature, humidity and wind speed where PV or a battery is present [1].

B. Edge Analytics Layer

The edge server stores each message in a local time-series database after identifying its source, then computes derived parameters from the cumulative counter kWh as in [2]:

$$E_{min} = kWh(t) - kWh(t - \Delta) \quad (1)$$

$$E_{today} = kWh(t) - kWh_{midnight} \quad (2)$$

$$E_{wk} = (1/7) \sum E_{md,k} \quad (k = 1 \dots 7) \quad (3)$$

$$D = (E_{today} - E_{wk}) / E_{wk} \times 100 \quad (4)$$

where $E_{md,k}$ is the maximum daily value on day k . In (4) the sign convention is chosen so that positive values mean above-average use; [2] states the percentage with the subtraction reversed but displays the same UP/DOWN semantics. If the midnight reading is missing because of an outage, the last available reading in a look-back window is used, which is the interval-based rule of [2]. Fleet totals are sums over meters. Electrical quality indicators follow [3] and [5]:

$$S = \sqrt{P^2 + Q^2}, \quad PF = P / S \quad (5)$$

$$LF = P_{avg} / P_{peak}, \quad PAR = P_{peak} / P_{avg} \quad (6)$$

Raw values stay at the edge. Only a filtered subset of key-value pairs is forwarded to the cloud, which cuts cloud storage and exposure [2]. A day-ahead LSTM forecast of PV output, load and battery state of charge runs at the edge, using the inputs and structure of [1]. An inactivity alarm is raised when a meter sends no data for a configurable period (ten minutes in [2]) and clears when data resume.

C. Control Layer

Control is automatic rather than dependent on user response [5]. Let P_{th} be a peak threshold. If $P(t) > P_{th}$ for a dwell time T_d , non-critical loads are shed in ascending priority order until $P(t) \leq P_{th}$; they are restored after a hysteresis margin. Critical loads are never shed. For an inductive load with active power P and lagging angle ϕ_1 , the capacitor bank supplies reactive power to reach a target angle ϕ_2 :

$$Q_c = P(\tan \phi_1 - \tan \phi_2) \quad (7)$$

Because (5) separates P , Q and S , the effect of compensation (lower S and current) can be reported separately from the effect of shedding (lower P), which [5] argues earlier studies blur.

D. Cloud Layer

The cloud provides remote dashboards, long-term storage of filtered data, policy updates and global analytics. Large historical storage and rule-engine execution are cheaper at the edge: [2] estimates about 529.7 rule-engine executions per minute for 19 devices, close to the limit of the cloud platform's minimum subscription plan.

E. Smart-Meter Signal Processing

At the meter, RMS values and powers are computed from sampled signals. In the design of [3], voltage and current are sampled at 4096 Hz and the RMS values over N samples are

$$V_{rms} = K_v (\sqrt{\sum v^2(n) / N} - v_{off}) \quad (8)$$

$$I_{rms} = K_i (\sqrt{\sum i^2(n) / N} - i_{off}) \quad (9)$$

where K_v and K_i are scaling factors and the offsets remove converter noise. Active power is the scaled mean of $v(n) \cdot i(n)$. Reactive power uses the voltage sample shifted by 90° , $v_{90}(n)$, in place of $v(n)$; apparent power follows from (5). Three-phase totals are the sums over phases. Accuracy depends on the sensing hardware: the industrial meters in [2] meet $\pm 1\%$ under IEC 60044-1 [8] for the wired meter and class 0.5S under IEC 62053-22 [9] for the wireless meter.

F. Edge Processing Procedure

The edge executes a fixed procedure on every incoming message. It combines the rule-chain structure of [2] (identify source, store, derive parameters, filter, alarm, report) with the control and security checks of Sections III-C and IV.

G. Worked Example: Power Factor Correction

The following illustrative calculation applies (5) and (7) to the induction-motor load of [5] (1.5 kW, 0.8 power factor lagging). It is not a measured result. Before correction, $S = 1.5/0.8 = 1.875$ kVA and $Q = \sqrt{(1.875^2 - 1.5^2)} = 1.125$ kvar. To reach $PF = 0.95$ ($\tan \phi_2 = 0.329$), (7) gives $Q_c = 1.5 \times (0.75 - 0.329) \approx 0.63$ kvar. After correction $S = 1.5/0.95 \approx 1.58$ kVA, a reduction of about 16% in apparent power and a similar fall in line current, while P is unchanged. This matches the qualitative finding of [5]: compensation lowers apparent power and current but leaves real energy use almost unchanged. The 41.58% reduction reported in [5] arises under that study's own loading and should not be read from this example.

IV. SECURITY MODEL

A. Threats

Three attack classes matter most. First, data attacks: false data injection and data poisoning of weather or meter inputs, and gradient-based adversarial perturbation of a trained forecaster. FGSM [6] perturbs an input x in the direction that increases the loss J :

$$x_{adv} = x + \epsilon \cdot \text{sign}(\nabla_x J(\theta, x, y)) \quad (10)$$

BIM [10] applies this step iteratively with clipping, and DeepFool [11] seeks a minimal perturbation across the decision boundary [1]. Second, communication attacks: eavesdropping on consumption patterns, which reveals occupancy, and access to unencrypted or weakly authenticated channels such as the MQTT setup in [2]. Third, availability attacks: denial of service, heartbeat spoofing or jamming, and fake base stations where 5G or multi-access edge computing is used [4].

B. Countermeasures

Adversarial training [6] minimises the worst-case loss within an ϵ -ball:

$$\min_{\theta} \mathbb{E} [\max_{\|\delta\|_{\infty} \leq \epsilon} L(\theta, x + \delta, y)] \quad (11)$$

Defensive distillation [7] trains a second network on the soft labels of a teacher whose softmax uses temperature T :

$$F(x) = \text{softmax}(Z(x) / T) \quad (12)$$

Heartbeat monitoring [4] declares a failure when no heartbeat arrives within T_f :

$$F(t) = 1 \text{ if } H(t) = 0 \text{ for } t > T_f, \text{ else } 0 \quad (13)$$

Following [1], two observation layers run together: the day-ahead forecast and an online check of measured values at the devices, so each validates the other. Countermeasures marked “proposed” are inferred from the gaps identified in [2] and [4] and have not been evaluated in those studies.

C. Layered Defence Design

The countermeasures are arranged in layers so that one failure does not defeat the whole system. At the device layer, meters authenticate to the broker and sign or encrypt payloads. At the network layer, TLS-protected MQTT, VPN isolation and a firewall limit who can reach the edge [2]. At the edge, input validation rejects physically impossible values before they reach the forecaster, and the two observation layers of [1] compare predicted and measured values. At the model layer, adversarial training and distillation harden the forecaster. At the orchestration layer, heartbeat and inactivity monitoring detect silent failure, and an operator is alerted when automatic actions fire repeatedly.

[4] adds design principles that fit this arrangement: modular separation of edge, fog and cloud layers; prioritised communication through 5G network slicing for control and heartbeat traffic; edge intelligence for fast local response;

continuous health monitoring; security by design; interoperability with open standards; and reproducibility through shared datasets and code. Where 5G is used, slices isolate critical control traffic from bulk telemetry, although [4] notes that fake base stations and misconfigured edge nodes remain open risks.

D. Control-Path Attacks

The reviewed studies attack either forecasts [1] or communication [2], but the framework also has a control path. An attacker who forges power readings could trigger shedding of non-critical loads without need, or hold readings below the threshold so that shedding never occurs. Three safeguards follow from the layers above: cross-checking each meter against the sum of its downstream meters or the main incomer; rate-limiting automatic actions with operator alerts; and keeping critical loads outside the automatic path altogether, as [5] does. These are proposals and are not evaluated in any reviewed study.

V. EVIDENCE REPORTED BY THE SOURCE STUDIES

A. Headline Results

The headline results below are as reported by the original authors. These are not results of this paper, and the studies differ in setting, scale and metric, so they give context rather than a benchmark.

Three observations follow. First, the results are complementary. [5] shows that reactive power compensation and load shedding act on different quantities, so a framework should report both. [2] shows that nine useful indicators come from one cumulative counter at about 5.6% CPU, so the edge layer is not the bottleneck at this scale. [3] shows that weekday and seasonal patterns are visible from a few meters, which supports setting thresholds from data.

Second, the security results need a careful reading. The conclusion of [1] states that accuracy does not fall below 80% under attack once defences are applied. Its own tables for the second model show lower values in some months: with distillation, April reaches 44.03% under FGSM and February 61.45%. The defences therefore help consistently but not uniformly, and the weakest months are those where the clean model was already weakest (April and February clean accuracy 79.25% and 76.80%). A deployment should monitor per-period accuracy rather than assume a fixed floor.

Third, no study evaluates control and security together. The attacks in [1] degrade a forecast; the loads in [5] are shed on measured values. An attacker who forges measured power could cause unnecessary shedding or suppress it, and this path is untested.

B. Attack and Defence Detail

Representative months from the SIEMS tables [1] for the two LSTM models under the three attacks, each with defensive distillation (DD). Under FGSM, model 1 in February falls from 96.98% to 24.87%, a drop of 72.11 percentage points, and distillation raises it to 67.6%. Model 1 in June recovers to 88.81%. Model 2 shows weaker behaviour: in April the clean accuracy is only 79.25%, FGSM reduces it to 16.37%, and distillation restores 44.03%.

The same study reports false positive rate (FPR), precision and recall for the attacks. For model 1 in January, FPR is 0.49% without attack and 64.20% under FGSM, so attacks do not only lower accuracy; they flood the detector with false alarms. For model 2 in February the FPR is already 22.87% without any attack, again showing that months with poor clean performance are the most exposed. These measures are defined in Section VI-D.

C. Scalability and Cost of Edge Processing

Edge load can be reasoned about from the rule-chain structure of [2]. Each of N meters passes through a 25-node chain every minute, a 27-node chain every 10 min, and contributes to a 33-node fleet chain every 10 min and a 6-node chain every hour. For $N = 19$ this gives

$$R = 25N + 27N/10 + 33/10 + 6/60 \quad (14)$$

$= 475 + 51.3 + 3.4 \approx 529.7$ rule-engine executions per minute, which agrees with [2]. The per-meter term is about 27.9 executions per minute, so a linear extrapolation to 100 meters would give roughly 2,800 per minute. This is an extrapolation, not a measurement: [2] measured 5.58% CPU and 16.356% memory only at 19 devices, on a modest edge machine, and the platform it uses states a limit of 1000 devices that depends on machine type. Dashboard loading for a month of history was also noted to be slower at the edge than in the cloud when many widgets are shown [2]. Edge hardware should therefore be sized by test before expansion, and fog-level partitioning [4] is a natural next step for larger sites.

Storage is also local. The edge in [2] is provisioned to keep five years of history on a 1 TB drive, while the meters in [3] hold only two to three months and rely on the server for

long-term records. Holding raw data at the edge keeps recurring cloud charges low and gives researchers a dataset for later anomaly detection, which [2] lists as future work and [4] identifies as a shortage in the field.

VI. DISCUSSION AND FUTURE WORK

A. Limitations of the Reviewed Evidence

- Scale. Test loads in [5] are small (non-critical loads of about 0.063 kW, 4% of monitored load), the cloud side of [2] covers one meter, and [3] reports consumption rather than control.
- Transport security. Unencrypted MQTT with username–password authentication [2] is below the standard the security model in Section IV assumes.
- Heartbeat integrity. Heartbeat studies rarely consider spoofing or jamming [4].
- Outages. In [2], an edge outage of a day or more delays new parameters by up to seven days.
- Data quality. Security studies in [4] often use datasets unrelated to grid traffic, and [1] uses one site's data.

B. Validation Plan

A testbed would combine three-phase meters, an ESP32-class controller with relays and a capacitor bank, an edge server running a rule engine and database, and an LSTM forecaster. FGSM, BIM and DeepFool perturbations would be injected into forecast inputs, and forged power readings into the control path. Proposed metrics are forecast error before, during and after attack; detection latency; peak demand, load factor and PAR; power factor and apparent power; shedding errors under spoofed readings; and edge CPU and memory use. Further directions include federated learning across sites, MARL-based dispatch [4], and extension from kWh-only analytics to reactive-power analytics.

C. Deployment Considerations

Retrofit is feasible because the building blocks are standard: current transformers on existing panels [2], [3], relays and capacitor banks on selected branches [5], and an edge server on the plant LAN. Sites should classify loads with plant staff first, since the shedding rule is only safe if the critical set is correct; [5] bases classification on production continuity, safety and redundancy. Data use also needs governance. Consumption traces reveal occupancy and production schedules [4], so access to raw data should be restricted, and only filtered values should leave the site. Finally, interpretation matters. [2] and [3] both observe that weekly averages that include non-working days distort

comparisons, so thresholds and deviation indicators should be computed per day type.

D. Evaluation Metrics

Forecast quality is measured by root mean square error and mean absolute percentage error over M samples:

$$\text{RMSE} = \sqrt{(1/M) \sum (y_m - \hat{y}_m)^2} \quad (15)$$

$$\text{MAPE} = (100/M) \sum |y_m - \hat{y}_m| / |y_m| \quad (16)$$

Attack detection is measured by precision, recall and FPR, as in [1], computed from true and false positives and negatives (TP, FP, TN, FN):

$$\text{FPR} = \text{FP}/(\text{FP} + \text{TN}), \quad \text{Precision} = \text{TP}/(\text{TP} + \text{FP}), \quad \text{Recall} = \text{TP}/(\text{TP} + \text{FN}) \quad (17)$$

Detection latency is the time from attack onset to alarm. Control effectiveness is measured by peak demand, load factor and PAR (6), power factor, and the number of unnecessary or missed shedding events under spoofed readings. System cost is measured by edge CPU and memory use and rule-engine executions per minute (14).

VII. CONCLUSION

This paper combined five studies on IoT-based energy management into a layered, secure edge-assisted framework for industrial facilities. The framework joins fine-grained metering, edge analytics, automatic demand-side control with power factor correction, forecasting, and defences against data and communication attacks. The reported evidence shows that each building block works on its own: lower apparent and peak power, low edge overhead, and recoverable forecast accuracy under attack. It also shows that their interaction is untested, that transport security and heartbeat integrity are open, and that defended accuracy varies by period. The next step is the experimental validation outlined in Section VI-B.

Several practical lessons follow for industrial operators. Start with metering and edge analytics, because they are low-risk and give the data that thresholds and forecasts need. Add automatic control for non-critical loads only after load classification has been agreed. Encrypt and authenticate every link before connecting the edge to a wider network. Treat forecast models as attackable components and track their accuracy by period. Future work will build the testbed of Section VI-B, extend the analytics from kWh to reactive power, test control-path attacks, and explore federated learning across sites and MARL-based dispatch [4].

REFERENCES

- [1] P. Asef, R. Taheri, M. Shojafar, I. Mporas, and R. Tafazolli, "SIEMS: A secure intelligent energy management system for industrial IoT applications," *IEEE Trans. Ind. Informat.*, 2022, doi: 10.1109/TII.2022.3165890.
- [2] A. A. Mirani, A. Awasthi, N. O'Mahony, and J. Walsh, "Industrial IoT-based energy monitoring system: Using data processing at edge," *IoT*, vol. 5, no. 4, pp. 608–633, Sep. 2024, doi: 10.3390/iot5040027.
- [3] M. U. Saleem, M. R. Usman, and M. Shakir, "Design, implementation, and deployment of an IoT based smart energy management system," *IEEE Access*, vol. 9, pp. 59649–59664, 2021, doi: 10.1109/ACCESS.2021.3070960.
- [4] R. B. Younis, M. Alkasassbeh, and A. Aldweesh, "Advanced technologies to smart grid systems: Challenges and demands," *IEEE Access*, vol. 13, pp. 117732–117752, 2025, doi: 10.1109/ACCESS.2025.3586466.
- [5] E. Tweve, G. Philipo, B. Kichonge, and T. Kivevele, "IoT-based demand-side energy management: Enhancing peak hour efficiency through automated control of appliances," *PLoS ONE*, vol. 21, no. 5, p. e0347990, May 2026, doi: 10.1371/journal.pone.0347990.
- [6] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv:1412.6572*, 2014.
- [7] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *Proc. IEEE Symp. Security Privacy (SP)*, 2016, pp. 582–597.
- [8] Instrument Transformers Part 1: Current Transformers, IEC 60044-1, 1996.
- [9] Electricity Metering Equipment Particular Requirements Part 22: Static Meters for AC Active Energy (Classes 0.1S, 0.2S and 0.5S), IEC 62053-22, 2020.
- [10] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," 2016.
- [11] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "DeepFool: A simple and accurate method to fool deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 2574–2582.