International Journal of AI, Engineering and Management Studies (IJAIEMS) E-ISSN: 3139-7034
Volume 1 | Issue 2 | Jul–Dec 2026
Page. no : 115-143

OPEN ACCESS

Artificial Intelligence for Smart Renewable Energy Grids: A Structured Literature Review and Integrated Framework

RatibSikder, Abdur R. Sikder

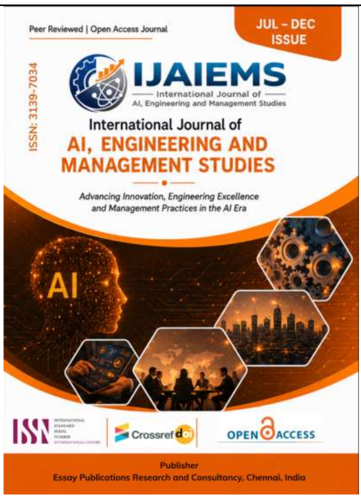

Article History

Received: 22th Aug 2026

Revised: 25th Aug 2026

Accepted: 28th Aug 2026

Published: 30th Aug 2026

DOI : 10.64771/ijaiems.v1.i2.10

Abstract:

Renewable energy is being added to the grid faster than the grid itself was ever built to handle. Solar and wind depend on the weather, so their output goes up and down in ways that traditional, mostly one-directional grid infrastructure was not designed to absorb. That mismatch can create voltage instability, energy imbalance, and inefficient dispatch [1], [2]. This paper is a structured literature review of how Artificial Intelligence (AI), including classical machine learning, deep learning, graph neural networks, physics-informed learning, explainable AI, and reinforcement learning, can be built into smart grid systems to deal with these problems. Using twenty-three peer-reviewed and technical-report sources found through a systematic search process, this review pulls together forecasting models for renewable generation and demand, topology-aware and physically-constrained state estimation, explainability methods that help build control-room trust, and reinforcement-learning controllers used for battery scheduling and microgrid dispatch. From there, the paper proposes an integrated AI-powered smart grid framework, walks through a real-world case study of duck-curve management at the California Independent System Operator (CAISO), and discusses the trade-offs, deployment challenges, and open research questions tied to each method. The evidence shows that AI-based forecasting and dispatch improve renewable utilization and supply-demand balance compared to conventional rule-based control, though results are context-specific. Cybersecurity, model transferability, data privacy, and computational scalability are still the main things standing in the way of wider deployment.

Keywords: *Smart grid; renewable energy; artificial intelligence; machine learning; deep learning; graph neural networks; physics-informed neural networks; explainable AI; reinforcement learning; literature review.*

HOW TO CITE

Artificial Intelligence for Smart Renewable Energy Grids: A Structured Literature Review and Integrated Framework. (2026). *International Journal of AI, Engineering and Management Studies (IJAIEMS)*, *1*(2), 115-143. https://doi.org/10.64771/ijaiems.v1.i2.10

Abbreviations

Abbreviation	Definition
AI	Artificial Intelligence
AMI	Advanced Metering Infrastructure
CAISO	California Independent System Operator
CNN	Convolutional Neural Network
DL	Deep Learning
DQN	Deep Q-Network
FDI	False Data Injection (attack)
GNN	Graph Neural Network
GRU	Gated Recurrent Unit
IoT	Internet of Things
LSTM	Long Short-Term Memory (network)
MAE / RMSE / MAPE	Mean Absolute Error / Root Mean Square Error / Mean Absolute Percentage Error
ML	Machine Learning
PINN	Physics-Informed Neural Network
PMU	Phasor Measurement Unit
PPO	Proximal Policy Optimization
RES	Renewable Energy Source(s)
RL	Reinforcement Learning
SCADA	Supervisory Control and Data Acquisition
XAI	Explainable Artificial Intelligence

1. INTRODUCTION

1.1 Background and Motivation

Climate change caused by human greenhouse gas emissions is one of the biggest engineering and policy problems of this century. The electricity sector has long depended on burning fossil fuels, and it remains one of the largest sources of global carbon emissions, which is a big part of why there has been such a strong international push to decarbonize power generation [1]. Over roughly the last ten years, that push has led to a fast build-out of solar and wind power capacity. Falling technology costs, supportive policies, and growing demand for clean electricity from both the public and corporations have all played a role. Renewable energy is now seen as central to both energy security and environmental sustainability, since it lowers dependence on imported fossil fuels while also cutting the grid's operational carbon footprint [1], [2].

Even so, solar and wind generation behave very differently than the dispatchable thermal power plants the grid was originally designed around. Solar output changes with cloud cover, time of day, and season. Wind output depends on atmospheric conditions that can shift within a few hours. Because of this intermittency, the amount of power being generated at any given moment often does not match the amount being demanded, so the grid has to absorb short-term imbalances through storage, flexible generation, or demand-side adjustments [1], [2]. Conventional power grids were built around centralized, dispatchable generation and a mostly passive, one-directional flow of electricity to consumers. That kind of architecture just was not designed for the two-way power flows, distributed generation, and rapid swings that come with high levels of renewable penetration [3].

Smart grids are the engineering response to that mismatch. By adding communication infrastructure, automation, and sensing on top of the physical power network, smart grids allow electricity and information to flow in both directions. That lets distributed solar and wind installations, battery storage systems, and even individual consumers take an active role in grid operation [1], [3]. Advanced Metering Infrastructure (AMI), Supervisory Control and Data Acquisition (SCADA) systems, and distributed control platforms give operators a level of visibility and control that older infrastructure simply could not provide [1], [13]. The problem is that the amount, speed, and variety of data these systems produce, weather forecasts, sensor telemetry, consumption patterns, is more than conventional rule-based control logic can realistically process in real time.

1.2 The Case for Artificial Intelligence in Grid Operations

This is the gap Artificial Intelligence is positioned to fill. Machine learning and deep learning models are good at pulling patterns out of large, messy, and varied datasets, which makes them useful for forecasting renewable generation and electricity demand [2], [4], [5]. Models such as Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNNs), and Transformer-based architectures have performed well in short-term load and generation forecasting, often doing better than classical statistical methods once there is enough training data [2], [5], [6]. Beyond modeling how things change over time, Graph Neural Networks (GNNs) explicitly capture the spatial layout and physical topology of power distribution networks. That gives them a promising way to do multi-node state estimation that actually respects how the grid is physically connected [19].

Physics-Informed Neural Networks (PINNs) take this further by embedding power-flow physics directly into how the network is trained, so the resulting state estimates stay consistent with basic circuit laws instead of just being statistical guesses [20]. Explainable AI (XAI) techniques give interpretable attributions for deep learning predictions, and this is increasingly treated as a requirement, not just a nice extra, for operators to trust control-room decisions [8], [21]. Reinforcement learning (RL) goes a step beyond forecasting and state estimation by enabling adaptive, sequential decision-making for problems like battery charge and discharge scheduling and demand response, where the best action depends on a system state that is always changing and never fully known [7], [9].

The scale of the renewable transition is what makes all of this feel urgent. Global renewable power capacity additions grew by 22% to reach almost 685 GW in 2024, a record high, and the International Energy Agency expects annual additions to keep climbing, driven mostly by solar PV and wind [11]. Rising renewable penetration combined with continued growth in electricity demand puts more pressure on grid operators every year. That strengthens the case for AI-assisted forecasting, state estimation, and control as practical, near-term tools that work alongside longer-term investments in transmission and storage infrastructure, not instead of them. AI-enabled monitoring also supports predictive maintenance and early fault detection, which matters a lot in places where distributed renewable assets and aging grid infrastructure exist side by side [15], [16].

1.3 Contributions of This Review

This paper makes four main contributions.

1. A structured review of the main AI methods used in smart renewable energy grids, covering classical machine learning, deep sequential models, graph neural networks, physics-informed learning, explainable AI, and reinforcement learning, organized by method type and drawn from twenty-three peer-reviewed and technical-report sources.

2. An integrated, engineering-level framework connecting weather and IoT data acquisition, AI-based forecasting, topology-aware state estimation, explainability, and RL-driven storage dispatch, shown both as an architectural diagram and as a formal control-loop algorithm.

3. A real-world case study, California ISO's management of its solar-driven "duck curve," that grounds the proposed framework's forecasting-to-storage-dispatch logic in an actual operating grid rather than only in simulated or laboratory settings.

4. A critical look at practical applications, quantitative benefits, deployment barriers, and open research directions, including cybersecurity, model bias and transferability, data privacy, regulatory acceptance, and computational scalability.

1.4 Paper Organization

The rest of this paper is organized as follows. Section 2 explains how the literature reviewed here was found and selected. Section 3 covers the smart grid and AI literature and lays out the research gap this paper is trying to address. Section 4 is a technical deep dive into the algorithmic families involved. Section 5 presents the proposed integrated framework and dispatch algorithm. Section 6 covers applications, and Section 7 walks through a real-world case study. Section 8 quantifies the benefits, Section 9 works through the challenges and security considerations, and Section 10 looks at future research directions. Section 11 discusses trade-offs across the reviewed methods, Section 12 is upfront about this review's limitations, Section 13 concludes, and Section 14 documents the evaluation metrics used throughout.

2. REVIEW METHODOLOGY

To keep the basis for this review transparent, this section documents the search strategy, screening criteria, and selection process used to find the studies discussed in Section 3. It follows the general spirit of structured review reporting, adapted to fit the scope of a conference-length paper.

2.1 Search Strategy and Databases

Candidate literature was found through keyword searches of IEEE Xplore, ScienceDirect, SpringerLink, MDPI, arXiv, and the technical-report repositories of the International Energy Agency (IEA) and the U.S. National Renewable Energy Laboratory (NREL), covering publications from 2010 through 2025. That range was chosen so it would include Farhangi's foundational 2010 description of the smart grid paradigm [3] along with the most recent (2024 to 2025) work on physics-informed state estimation and renewable capacity forecasting [11], [20]. Search terms were combined using Boolean operators across five thematic clusters that match the paper's technical scope.

- Infrastructure and grid architecture: "smart grid" AND ("Advanced Metering Infrastructure" OR "SCADA" OR "phasor measurement unit").
- Forecasting: ("renewable energy" OR "load") AND "forecasting" AND ("machine learning" OR "deep learning" OR "LSTM" OR "transformer").
- Spatial and physics-constrained modeling: "graph neural network" AND "power grid"; "physics-informed neural network" AND "power system".

- Control, explainability, and security: "reinforcement learning" AND ("microgrid" OR "energy storage"); "explainable AI" AND "energy"; "false data injection" AND "smart grid".
- Emerging directions: "federated learning" AND "smart grid"; "digital twin" AND "power system"; "peer-to-peer energy trading".

2.2 Screening and Selection Funnel

An initial keyword search across the clusters above turned up about 210 candidate records after obvious duplicates across databases were removed. Screening titles and abstracts against the inclusion criteria in Section 2.3 narrowed this down to 61 records considered potentially relevant. Full-text screening of those 61 records, checking that each one reported a distinguishable method, at least one qualitative or quantitative finding, and enough methodological detail to summarize its dataset, technique, and limitations, left the twenty-three sources used in Section 3. This funnel (about 210 identified, 61 screened at title and abstract level, 23 selected after full-text review) is reported here for the sake of transparency. As Section 12 points out, it reflects a targeted, single-reviewer screening process rather than the dual-reviewer, formally registered protocol used in a fully systematic review.

2.3 Inclusion and Exclusion Criteria

Studies were included if they (a) applied a machine learning, deep learning, graph-based, physics-informed, reinforcement-learning, or explainable-AI method to a smart grid, microgrid, or renewable-integration problem, (b) reported a distinguishable method along with at least one qualitative or quantitative finding, and (c) gave enough detail to summarize their dataset, method, application, and reported limitations. Studies were excluded if they were purely promotional or vendor material with no described methodology, if they were duplicate reports of the same underlying study, or if they fell outside the 2010 to 2025 window without being foundational to the field. Farhangi's 2010 paper [3] was kept as a deliberate exception, since it still plays a foundational role in smart grid architecture discussions.

2.4 Note on Reference Verification

A couple of citations are worth a quick note. Alsaigh, Mehmood, and Katib's review [8] exists as both a 2022 arXiv preprint and a 2023 peer-reviewed Frontiers in Energy Research paper; the peer-reviewed version is the one cited here. Nagappan et al.'s review [9] carries an article identifier (2023ss070) suggesting it was accepted in 2023, but its formal volume publication is dated 2024; this paper follows the 2024 volume-publication year used in Table 1 and the reference list, though some other citing sources still list it as 2023. A couple of other sources ([5], [10]) come from an expanded abstract accepted before this full review was written, and they're cited with the bibliographic detail available at that time.

3. LITERATURE REVIEW

Before getting into individual studies, Fig. 1 maps out how the topical clusters reviewed in this section connect along the forecasting-to-dispatch pipeline that motivates the pipeline proposed in Section 5.

Fig. 1. Literature map: how the topical clusters reviewed in Section 3 connect along the forecasting-to-dispatch pipeline (author-generated; section references in parentheses).

3.1 Smart Grid Technology & Infrastructure

A smart grid extends the conventional power network with communication, automation, and sensing layers that let electricity and information flow in both directions [1], [3]. Farhangi's foundational description of the smart grid frames it as a system-level integration of self-healing capability, active consumer participation, and support for many different generation and storage options [3]. Advanced Metering Infrastructure (AMI) sits at the center of that integration [1], [13]. SCADA systems and distributed control platforms give operators real-time monitoring and remote control, and phasor measurement units (PMUs) improve situational awareness across the network [1].

Depuru, Wang, and Devabhaktuni's early survey of smart meter deployment lays out the challenges, issues, and status of AMI rollouts. One thing they point out is that the same communication channels that enable remote metering and demand-side visibility also widen the attack surface available to adversaries [13]. That tension between openness and exposure keeps showing up throughout the literature, and it is a big part of why Section 9 spends so much time on cybersecurity. Kiasari, Ghaffari, and Aly synthesize how AMI, SCADA, and machine learning work together to support renewable energy source (RES) integration at scale. They frame AI not as a replacement for this infrastructure but as an analytical layer that sits on top of it [1].

3.2 Machine Learning & Deep Learning Architectures

Artificial Intelligence provides the analytical power needed to turn raw smart grid data into forecasts and control decisions that operators can actually act on [2], [4]. Classical machine learning methods, things like Support Vector Machines, Random Forest, and XGBoost, can model nonlinear relationships without needing huge training sets [2]. Deep learning methods (LSTM, GRU, Transformers) learn hierarchical representations straight from high-resolution time-series data [2], [5]. Biswal, Deb, and colleagues report that across a range of smart grid load datasets, deep learning approaches tend to beat statistical baselines on standard error metrics like MAPE and RMSE, though how much better appears to depend a lot on the dataset [2].

Kong et al. show that LSTM-based architectures improve demand-prediction accuracy over simpler baseline models on smart grid demand data. Their evaluation only covers a single region, though, so how well it would generalize across regions is still an open question [5]. Barhmi et al. review several machine learning models on solar and wind meteorological data and find that how well these models forecast really comes down to the quality of the underlying weather-sensor inputs [6]. One pattern shows up across all these studies pretty clearly: simpler models like Random Forest hold up fine at very short horizons where nonlinear relationships are limited, but sequence-aware deep architectures pull ahead once the horizon gets longer and long-range temporal dependencies start to matter more for prediction error [2].

3.3 Spatial Topology: Graph Neural Networks (GNNs)

Classical ML and temporal deep architectures mostly treat grid measurements as isolated time-series streams. But modern power distribution systems actually have physical graph topologies, and Graph Neural Networks (GNNs) capture spatial node connections (substations, transformers, smart meters) by representing the network as a graph and passing information along edges that correspond to real electrical connections. Kundacina, Cosovic, and Vukobratovic show that a GNN trained on synthetic phasor-measurement-unit (PMU) datasets can learn to predict power-system state estimates, complex bus voltages, directly from PMU voltage and current measurements. They report accurate predictions across multiple test scenarios and also look at how sensitive the model is to missing input data [19]. This line of work helps explain why topology-aware learning seems to matter in practice. A fault or voltage excursion at one node propagates along edges defined by real line impedances, not arbitrary feature correlations, and a GNN's message-passing mechanism is built specifically to try to model that kind of propagation.

3.4 Physics-Informed Neural Networks (PINNs) in Power Systems

Pure data-driven deep learning models can hit high statistical accuracy, but they run the risk of producing non-physical states under conditions or contingencies they have not seen before. Physics-Informed Neural Networks (PINNs) get at this problem by embedding governing physical equations directly into the training loss function, so the network gets penalized for predictions that violate those equations, not just for predictions that stray from the labeled data. Ngo, Nguyen, Vu, Zhang, and Ngo apply this idea to power system state estimation with a physics-informed graphical neural network, combining the topology-awareness of graph learning with power-flow physics constraints so the estimated states end up more physically consistent [20]. Here is the basic idea: a purely data-driven model might reach a low training loss on the statistical objective alone and still produce voltage or current estimates that break power-flow physics under contingency conditions it never saw during training. The physics-constrained loss term in a PINN directly penalizes those kinds of violations, trading a slightly slower or higher training loss for state estimates that are more physically plausible.

3.5 Explainable AI (XAI) for Control Room Integration

A major bottleneck in getting complex deep neural networks and reinforcement learning policies into real control rooms is the black-box nature of neural predictions [8]. Alsaigh, Mehmood, and Katib frame this as a governance problem just as much as a technical one. They argue that AI explainability and accountability frameworks are a prerequisite, not an afterthought, for regulators to accept AI-driven control actions in energy systems [8]. To build that kind of human-in-the-loop trust, Explainable AI (XAI) frameworks such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) are being built into prediction pipelines. Utama, Meske, Schneider, Schlatmann, and Ulbrich apply and compare several XAI tools for photovoltaic fault detection, finding that feature-attribution outputs help engineers see which measured quantities are driving a given fault classification. That does come with some extra computational overhead compared to a model with no explanation attached [21]. Notably, [8] and [21] address different aspects of explainability: [8] focuses on governance frameworks, while [21] evaluates specific XAI tools for fault diagnosis. A grid could have the governance framework Alsaigh et al. call for and still lack the kind of tool-level validation Utama et al. provide, or vice versa, so the two papers are more complementary than overlapping. These techniques give operators feature attributions for forecasting and diagnostic outputs, so they can check whether a model's output is being driven by legitimate meteorological or electrical signals rather than by a sensor glitch.

3.6 Reinforcement Learning & Energy Management

Reinforcement learning (RL) treats battery scheduling, energy dispatch, and demand-response as sequential decision-making problems [7], [9]. RL agents learn control policies by interacting with the grid environment and getting feedback in the form of reward signals. Albogamy et al. show, in a simulated grid environment, that an RL policy improves dispatch efficiency over rule-based control. The reported gains do come with a simulation-to-real gap that has not been closed yet [7]. Nagappan et al. provide a systematic review of AI approaches for supporting diverse demands in distributed smart grids, looking at how AI can help improve renewable energy integration, power management, and overall grid social welfare [9]. The sim-to-real gap is addressed differently in [7] and [18]. Albogamy et al. [7] acknowledge an unresolved gap in stationary battery dispatch, while Abdullah, Gastli, and Ben-Brahim's review of RL-based electric-vehicle charging management [18] reports field-adjacent data showing that learned scheduling policies reduce peak-load impact and charging cost compared to static time-of-use schemes by adapting in real time to occupancy, price, and grid-constraint signals, suggesting that RL deployment readiness varies by application. That difference is a reminder that "reinforcement learning works for grid control" is not one uniform claim. How close it is to real deployment depends a lot on which specific application is being discussed.

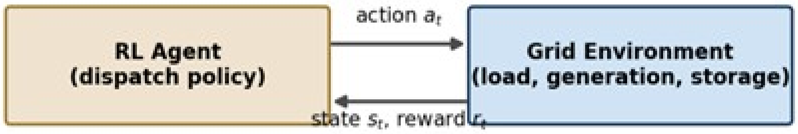

Fig. 2. Reinforcement-learning control loop for battery / storage dispatch (author-generated schematic, conceptual illustration synthesized from [7], [9], [18]).

3.7 Summary of Reviewed Literature

Table 1 pulls together the reviewed sources above, organized by dataset context, AI method, application, key findings, and reported limitations.

TABLE 1 Comparison of Reviewed Studies

Author(s) & Year	Dataset / Context	AI Method	Application	Key Findings	Limitations
Kiasari, Ghaffari, & Aly [1] (2024)	Smart grid literature corpus	Review / ML-supported architecture	RES integration, AMI, SCADA	Synthesizes AMI/SCADA/ML roles in RES integration	Review-level; no new experimental data
Biswal, Deb, et al. [2] (2024)	Smart grid load datasets (review)	ML & DL (SVM, RF, DNN, CNN, LSTM)	Short-term load forecasting	DL/ML outperform statistical baselines on MAPE/RMSE	Performance varies by dataset; limited cross-study comparability
Farhangi [3] (2010)	Conceptual grid architecture	Architectural framework	Smart grid paradigm design	Establishes self-healing, consumer-interactive design	Predates modern deep learning methods
Bishaw et al. [4] (2024)	Renewable energy systems survey	ML/DL applications survey	Forecasting, optimization, XAI	Surveys ML/DL/XAI roles in solar & wind RE system integration	Survey scope; accuracy not standardized across studies
Kong et al. [5] (2017)	Smart grid demand data	Deep learning (LSTM-based)	Demand prediction	Improved prediction accuracy over baselines	Single-region dataset; cross-domain generalization untested
Barhmi et al. [6] (2024)	Solar/wind meteorological data	ML forecasting models	Solar & wind forecasting	Reviews forecasting techniques; accuracy dependent on sensor data quality	Review-level; no new experimental data
Albogamy et al. [7] (2022)	Simulated grid	Reinforcement learning	Energy managemen	RL policy improves dispatch efficiency	Simulation-based; sim-to-

Author(s) & Year	Dataset / Context	AI Method	Application	Key Findings	Limitations
	environment		t & dispatch		real gap unresolved
Alsaigh et al. [8] (2023)	Smart energy governance framework	XAI & governance framework	AI explainability & governance	Identifies governance/explain ability needs	Conceptual; no algorithmic benchmarking
Nagappan et al. [9] (2024)	Smart grid literature	AI powered review	Renewable integration and energy managemen t	Reviews AI approaches for smart grid applications	Review-level; no new experimental data
AsadiAghajar i et al. [10] (2025)	Load & solar generation data	ML forecasting techniques	Load & solar forecasting	Improved short-term forecasting accuracy	Journal-review level
IEA [11] (2025)	Global capacity & market data	Statistical / forecast report	Renewable capacity outlook	685 GW added in 2024; ~95% solar+wind through 2030	Forecast; sensitive to policy revisions
Jafari et al. [12] (2023)	Smart grid / smart city literature	Review	Digital twin technology	Surveys DT applications in monitoring, fault diagnosis	Review-level; interoperability barriers noted
Depuru, Wang, &Devabhakt uni [13] (2011)	Smart metering infrastructur e	Survey & analytical framework	Smart meter security & AMI	Identifies AMI vulnerabilities and status	Historical perspective (pre-modern-DL era)
Fekri, Grolinger, & Mir [14] (2022)	Distributed smart meter data	Federated learning + RNN	Distributed load forecasting	Matches centralized-model accuracy while preserving privacy	Vulnerable to non-IID data across nodes
Mirzaee et al. [15] (2022)	Smart grid security literature	Review	Smart grid security & privacy	Surveys FDI/DoS threats and ML-based defenses	Review-level; evolving threat landscape
Ruan et al. [16] (2023)	Smart grid cybersecurity literature	Review (deep learning)	Cybersecurit y in smart grids	DL improves anomaly/intrusion detection over rule-based methods	Review-level; detection models themselves attackable

Author(s) & Year	Dataset / Context	AI Method	Application	Key Findings	Limitations
Cole, Ramasamy, &Turan [17] (NREL, 2025)	U.S. battery cost literature	Bottom-up cost model	Utility-scale storage economics	2024 storage cost ~$334/kWh, declining toward 2035	Cost projections; scenario-dependent
Abdullah, Gastli, & Ben-Brahim [18] (2021)	EV charging management literature	Review (RL-based)	RL-based EV charging management	Learned policies reduce peak-load impact & cost	Review-level; results vary by study
Kundacina, Cosovic, &Vukobratovic [19] (2022)	Synthetic PMU datasets	Graph Neural Network	State estimation	Accurate state predictions; robustness to missing data analyzed	Synthetic data; sensitivity to topology changes
Ngo et al. [20] (2024)	Power system state estimation	Physics-informed GNN	Dynamic state estimation	Embeds power-flow physics; improves physical consistency	Requires accurate network parameters
Utama et al. [21] (2023)	Photovoltaic fault dataset	XAI (SHAP/LIME comparison)	PV fault detection	Feature attribution aids fault diagnosis	Added computational overhead
Sakhnini et al. [22] (2021)	Cyber-physical grid dataset	Ensemble deep learning	Physical-layer attack detection	High detection/localization accuracy for tampering	Complexity grows with network size
Shrestha et al. [23] (2019)	Nepal community microgrid case study	Review & case study	Peer-to-peer energy trading	P2P trading increases local renewable self-consumption	Case-specific; scalability constraints

3.8 Synthesis of the Reviewed Literature

Examining the studies in Table 1 side by side reveals a fairly coherent progression, not just a random pile of unrelated solutions. Early foundational work established the smart grid as an architectural concept and flagged AMI-level security concerns [3], [13]. A second cluster of studies applied classical and deep learning methods to the forecasting problem, moving gradually from statistical baselines toward LSTM and GRU architectures as more training data and compute became available [2], [5]. A third cluster realized that forecasting alone does not solve grid operation, and turned to graph-based and physics-constrained learning to respect the network's spatial and physical structure [19], [20]. A fourth cluster addressed the trust problem that resulted through explainability [8], [21], and a fifth applied reinforcement learning to the downstream control and dispatch problem [7], [9], [18]. Federated learning, digital twins, and dedicated cybersecurity studies show up as more recent additions, and they point toward the future-work directions covered in Section 10 [12], [14], [16].

One limitation shows up again and again across almost all of these studies: the gap between controlled evaluation (simulated environments, synthetic PMU datasets, single-region datasets) and how messy real-world deployment actually is. That gap is a big reason Section 5.7 spends time on computational complexity, and it is also why scalability and transferability get treated as first-class challenges in Section 9 rather than an afterthought.

3.9 Research Gap

Although the literature reviewed above covers renewable energy forecasting, topology-aware state estimation, explainability, and reinforcement-learning-based control, it mostly treats these as separate research threads. Forecasting studies typically evaluate model accuracy on its own, separate from downstream dispatch decisions [2], [5]. Graph-based and physics-informed state-estimation studies focus on estimation accuracy without connecting it to a control policy [19], [20]. Explainability studies evaluate attribution quality without embedding it into an actual operational pipeline [8], [21]. And reinforcement-learning studies for dispatch usually assume clean, pre-processed state information rather than deriving it from a topology-aware estimator [7], [9], [18]. Within the specific set of sources reviewed here, very few look at how all these techniques interact within a single, end-to-end pipeline that carries raw telemetry through forecasting, physically-consistent state estimation, explainable output, and reinforcement-learning-based dispatch as connected stages of one system. That is a claim about this review's 23-source sample, not a claim that no such integrated work exists anywhere in the broader literature.

That gap, the lack of an integrated treatment connecting forecasting, topology-aware and physics-constrained state estimation, explainability, and reinforcement-learning dispatch, is what motivates the conceptual framework proposed in Section 5. The framework is offered as a structured synthesis of how these complementary techniques could be combined, grounded in what each one has already shown it can do in the literature, rather than as a brand-new algorithm on its own.

4. ARTIFICIAL INTELLIGENCE TECHNOLOGIES: TECHNICAL DEEP DIVE

4.1 Classical Machine Learning

Random Forest and Decision Trees combine predictions from many decision trees trained on bootstrapped subsets of data. This reduces variance and keeps the model fairly interpretable [2]. Each tree splits the feature space (historical load, temperature, irradiance, calendar effects) into piecewise-constant regions, and averaging across the whole ensemble smooths out the sharp boundaries any single tree would produce on its own. XGBoost and other gradient-boosting frameworks build weak learners one at a time, with each one correcting the residual error left over from the ensemble so far. That gets high accuracy on structured tabular smart grid telemetry [2]. These classical methods share two practical advantages that keep them relevant even as deep learning has advanced: they train quickly on modest hardware, and their feature importances can be inspected directly without needing a separate XAI pipeline.

4.2 Deep Learning & Temporal Models

Convolutional Neural Networks (CNNs) pull local temporal features out of sliding windows of time-series data by applying learned filters across short segments of the sequence. In practice, they act like automated feature extractors for patterns such as daily load ramps or short irradiance dips [2]. Long Short-Term Memory (LSTM) networks use gated memory cells, an input gate, a forget gate, and an output gate, to hold onto long-range temporal dependencies. This avoids the vanishing-gradient problem that limits simpler recurrent networks in multi-step forecasting [5]. Gated Recurrent Units (GRUs) simplify the LSTM gating mechanism, which cuts down the parameter count for resource-constrained edge deployments while keeping most of the same long-range modeling capacity [2]. In practice, choosing between CNN, LSTM, and GRU architectures for a given smart grid forecasting task usually comes down less to which model is theoretically strongest and more to how much historical telemetry is available, how fast inference needs to run, and whether the deployment target is a cloud-scale forecasting service or an edge device with limited memory.

4.3 Spatial Topology & Physics-Informed Learning

Standard deep learning models treat data as unstructured arrays and ignore how physical power networks are actually laid out. Graph Neural Networks instead map substations, transformers, and meters as nodes, and transmission lines as edges. By running graph convolutions, repeatedly updating each node's representation based on its neighbors' representations, GNNs model spatial interactions across the feeder network. This is the mechanism Kundacina, Cosovic, and Vukobratovic apply directly to PMU-based state estimation [19]. Physics-Informed Neural Networks extend this idea by enforcing power-balance and bus-voltage physics constraints inside the training loss function itself, an approach Ngo et al. combine with graph-based learning to get state estimates that are both topology-aware and more physically consistent [20]. That combination shows up directly in the state-estimation stage of the proposed framework in Section 5, where a topology-aware, physics-constrained estimator produces the state information that the reinforcement-learning optimizer relies on downstream.

4.4 Reinforcement Learning & Energy Storage

Q-Learning is a foundational, value-based RL algorithm where an agent learns the expected long-term reward of taking a given action in a given state, gradually converging on an optimal control policy for problems like battery charge and discharge scheduling [7]. Deep Q-Networks (DQN) extend Q-Learning by approximating the action-value function with a neural network, which lets RL scale up to the larger, continuous state spaces you find in realistic grid environments. Proximal Policy Optimization (PPO), a policy-gradient method, directly learns a control policy while limiting how much each update can shift the policy, which improves training stability. It has been applied to multi-objective energy dispatch and demand-response problems where cost, reliability, and battery degradation all need to be balanced at the same time [7], [9]. Framing storage dispatch this way, as a sequential, reward-driven decision process rather than a single number to minimize, is what lets RL controllers anticipate generation and demand swings instead of just reacting to them after the fact [7], [9], [18].

4.5 Explainable AI & Feature Attribution

Explainable AI techniques like SHAP and LIME compute feature-contribution scores for neural network outputs. SHAP values come out of cooperative game theory. They split a prediction's deviation from the average output across the input features in a way that satisfies certain consistency and local-accuracy properties. LIME instead fits a small, locally interpretable surrogate model around one specific prediction. In grid control centers, XAI helps operators understand why an AI model forecasted a demand spike, flagged a fault, or recommended a particular battery action, which supports safer human-in-the-loop decision-making [8], [21].

4.6 Algorithmic Comparison

Table 2 lays out the advantages, limitations, and typical smart grid application of each algorithmic family discussed above. It gives one reference point for choosing a method based on a specific engineering constraint, like available compute, interpretability requirements, or topology awareness.

TABLE 2 Comparison of AI Algorithms: Advantages and Limitations

Algorithm	Advantages	Limitations	Typical Application
Random Forest / XGBoost	Robust to noise; interpretable; fast training	Limited long-range temporal modeling	Short-term load & tabular forecasting [2]
LSTM / GRU	Captures long-range sequential dependencies	Sequential training; higher compute cost than classical ML	Solar, wind & load forecasting [5]
Graph Neural Networks (GNN)	Explicitly models physical grid topology	Requires accurate topological maps; sensitive to topology changes	State estimation & spatial power flow [19]
Physics-Informed Neural Networks	Improves physical consistency of estimates; can reduce data need	Requires accurate network parameters	Dynamic state estimation [20]
Explainable AI (SHAP / LIME)	Provides feature attribution & transparency	Additional computational latency	Control-room operator support [8], [21]
Reinforcement Learning (DQN / PPO)	Adaptive, sequential decision-making	Reward-design difficulty; sim-to-real gap	Battery scheduling &microgrid dispatch [7], [9], [18]

Beyond the per-algorithm comparison in Table 2, it helps to place each method family on a common maturity scale, since methods at different stages of maturity carry different deployment risk even when their reported accuracy looks similar. Table 3 does this for the five method families most central to the proposed framework in Section 5.

TABLE 3 Maturity and Deployment-Readiness of Key Method Families

Method	Strength	Weakness	Maturity for Grid Deployment
LSTM / GRU	High forecasting accuracy on sequential data	Sequential training; limited spatial awareness	Mature — widely deployed in production forecasting [2], [5]
Transformer	Strong long-range dependency modeling	High compute & data requirements	Emerging — growing adoption, not yet standard [2]
Graph Neural Network (GNN)	Explicit grid-topology awareness	Implementation complexity; topology-change sensitivity	Emerging — strong research results, limited field deployment [19]
Physics-Informed Neural Network (PINN)	Physical consistency of estimates	Training complexity; needs accurate network parameters	Early adoption — promising but nascent in power systems [20]

Method	Strength	Weakness	Maturity for Grid Deployment
Multi-Agent Reinforcement Learning (MARL)	Decentralized, scalable control	Inter-agent coordination challenges; reward design	Research stage — largely confined to simulation studies [7], [9]

As Table 3 shows, the proposed framework in Section 5 deliberately combines mature components (LSTM/GRU forecasting) with emerging components (GNN state estimation, PINN physical constraints), rather than assuming every stage of the pipeline is equally ready for deployment. That distinction directly shapes the phased implementation discussion in Section 10.6.

4.7 Mathematical Formulations

To make the comparisons in Sections 4.1 through 4.6 more concrete, this subsection lays out the core mathematical formulation behind each method family, following the standard formulations used across the reviewed literature [2], [7], [19], [20].

LSTM Cell Update

An LSTM cell updates its hidden state h_t and cell state c_t at each timestep using three gates, forget (f_t), input (i_t), and output (o_t), each a sigmoid function of the current input x_t and previous hidden state h_{t-1}:

$$f_t = \sigma(W\_f \cdot [h_{t-1}, x_t] + b\_f),\ \ i_t = \sigma(W\_i \cdot [h_{t-1}, x_t] + b\_i),\ \ o_t = \sigma(W\_o \cdot [h_{t-1}, x_t] + b\_o) \quad (1)$$

The cell state is then updated as $c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W\_c \cdot [h_{t-1}, x_t] + b\_c)$, and the hidden state as $h_t = o_t \odot \tanh(c_t)$, where $\odot$ denotes elementwise multiplication. The forget gate f_t is what lets the network selectively keep or throw away information across long sequences, which directly addresses the vanishing-gradient limitation of simple recurrent networks discussed in Section 4.2 [2].

Graph Neural Network Message Passing

A GNN layer updates each node's representation by aggregating information from its neighbors, following the general message-passing formulation Kundacina et al. apply to power-system state estimation [19]:

$$h\_v^{(k+1)} = \varphi\big(h\_v^{(k)}, \square\_{u \in N(v)}\ \psi(h\_v^{(k)}, h\_u^{(k)}, e\_{vu}) \big) \quad (2)$$

where $h\_v^{(k)}$ is node v's representation at layer k, N(v) is the set of neighboring nodes (physically connected buses or meters), e_{vu} is the edge feature (for example, line impedance) connecting v and u, ψ is a learned message function, □ is a permutation-invariant aggregation such as summation or averaging, and φ is a learned update function. This is the mechanism by which a fault or voltage excursion at one node propagates through the model along the same edges that carry it physically through the network, as discussed in Section 3.3 and shown in Fig. 4.

Physics-Informed Loss Function

A PINN adds a physics-residual term to the standard supervised loss that penalizes violations of the governing power-flow equations, following the general formulation Ngo et al. apply to state estimation [20]:

$$L\_total = L\_data(\hat{y}, y) + \lambda \cdot L\_physics(\hat{y}),\ \ \text{where } L\_physics = \| g(\hat{y}) \|^2 \quad (3)$$

whereL_data is the standard supervised loss (for example, mean squared error against labeled measurements), $g(\hat{y})$ is the residual of the power-flow equations evaluated at the predicted state $\hat{y}$ (ideally zero for a physically consistent state), and λ is a weighting hyperparameter that balances data fidelity against physical consistency. As discussed in Section 3.4, this term is what lets a PINN get penalized for predictions that fit the training data but violate power-flow physics under conditions the training data never covered.

Reinforcement Learning: Bellman Optimality

The RL controllers discussed in Sections 3.6 and 4.4 learn a policy by approximating the optimal action-value function, which satisfies the Bellman optimality equation:

$$Q^*(s, a) = E[\, r + \gamma \cdot \max_{a'} Q^*(s', a') \mid s, a \,] \tag{4}$$

where $Q^*(s, a)$ is the optimal expected cumulative reward of taking action a (for example, a battery charge/discharge decision) in state s (for example, current net load and state of charge), r is the immediate reward received, $\gamma \in [0, 1]$ is a discount factor weighting future rewards against immediate ones, and s' is the resulting next state. DQN approximates Q^* with a neural network and trains it to minimize the temporal-difference error between the two sides of Equation (4). PPO instead directly parameterizes and optimizes a stochastic policy $\pi(a|s)$, subject to a constraint on how far each update can move the policy, which improves training stability compared to naive policy-gradient methods [7].

5. PROPOSED INTEGRATED AI SMART GRID FRAMEWORK

5.1 Architectural Overview

The framework proposed here is a conceptual synthesis grounded in the literature reviewed in Section 3, not an implemented or benchmarked system. It connects environmental and sensor telemetry to consumer energy delivery through a layered pipeline, shown in Fig. 3. A Weather Forecast Feed and IoT Sensor Network (smart meters, PMUs) stream raw telemetry into Cloud Ingestion and Data Processing modules. From there, data moves to a forecasting and state-estimation layer, where topology-aware, physics-constrained estimates are computed alongside generation and demand forecasts. An explainability layer then generates feature attributions for these outputs before they reach the reinforcement-learning optimization engine, which issues battery storage and distribution control actions to balance supply and demand dynamically.

5.2 Data Acquisition and Preprocessing Layer

The lowest layer of the framework brings together two very different data streams: weather forecast feeds providing irradiance, cloud cover, wind speed, and temperature projections, and an IoT sensor network of smart meters and PMUs providing electrical telemetry at sub-second to minute resolution. Before any of this data reaches a learning model, it passes through cleaning and sanitization steps that filter out sensor noise and screen for signs of False Data Injection (FDI) attacks. This matters because a forecasting or state-estimation model trained on corrupted telemetry can carry that corruption into every downstream decision [15], [22].

5.3 Forecasting and State Estimation Layer

This layer does two things at once. First, a topology-aware, physics-constrained estimator, combining Kundacina et al.'s graph-based approach with Ngo et al.'s physics-informed loss formulation, takes the cleaned telemetry along with the physical topology of the feeder network and produces a state estimate at every node [19], [20]. Second, a forecasting model (LSTM, GRU, or an ensemble, informed by the comparisons in Sections 4.1 and 4.2) projects near-term solar generation, wind generation, and demand [2], [5]. The gap between forecast demand and forecast generation is the net load that the downstream optimizer has to balance.

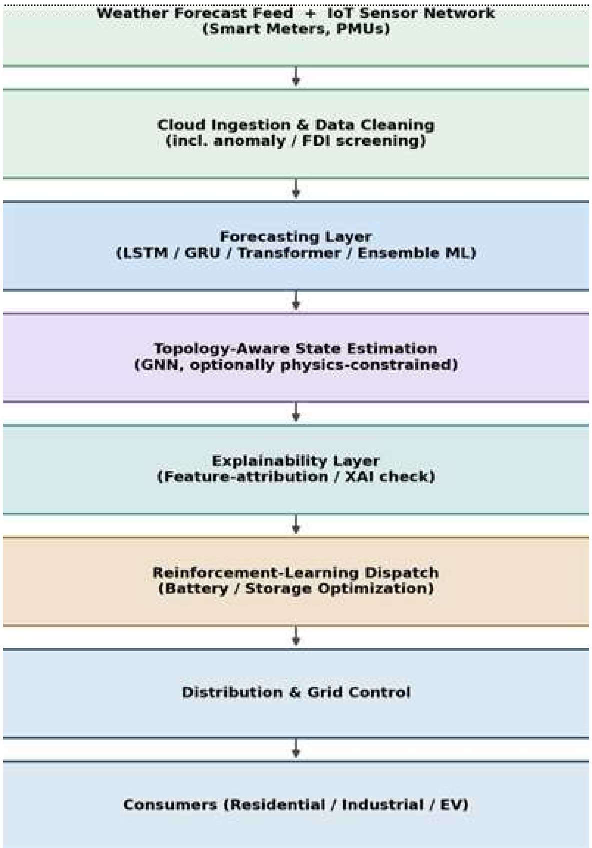

Fig. 3. Conceptual architecture of the proposed integrated AI smart-grid pipeline (author-generated diagram, synthesized from the reviewed literature; not a diagram of any single cited source).

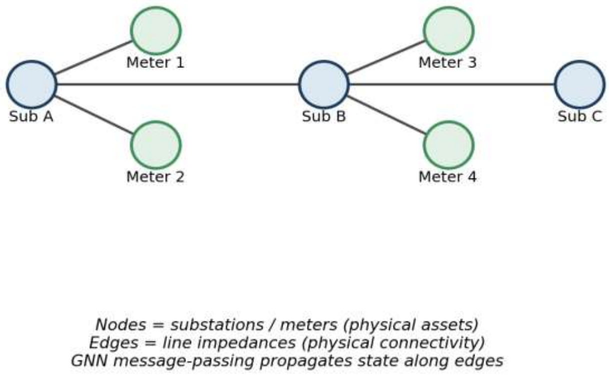

Fig. 4. Conceptual illustration of graph-structured grid topology used by GNN-based state estimation (author-generated schematic).

5.4 Explainability Layer

Before any forecast or state estimate reaches the reinforcement learning optimizer, an XAI module computes feature attributions for the forecast, showing which input features, solar irradiance, temperature, historical load, or time of day, are driving the current prediction [8], [21]. This step exists specifically to give operators a way to step in. If the attribution pattern does not match plausible meteorological or demand drivers, the operator can flag the input as potentially anomalous before it ever influences dispatch.

5.5 Reinforcement Learning Optimization Engine

The optimization engine takes in the net-load estimate, the state-estimation output, and the current battery state of charge, and computes a dispatch action using a DQN or PPO policy [7], [9]. That action specifies whether to charge, discharge, or hold each storage unit, and how to allocate distribution-level control actions across the feeder. Because the policy is trained against a multi-objective reward balancing cost, voltage stability, and battery degradation, its output is not simply the cheapest choice in isolation. It respects the physical and operational constraints surfaced by the earlier layers.

5.6 Algorithm 1: Forecasting, State Estimation, and Dispatch Cycle

Input: weather feed W, sensor stream S, topology graph G, battery state B

1: loop for each scheduling interval t

2: D_t ← CollectData(W, S) // weather + IoT telemetry

3: D_t ← CleanAndScreen(D_t) // noise filtering, FDI screening [15],[22]

4: X_t ← StateEstimate(G, D_t) // GNN + physics-informed [19],[20]

5: G_hat ← ForecastGeneration(D_t) // solar + wind (LSTM/GRU) [5]

6: L_hat ← ForecastDemand(D_t) // load forecast [2],[5]

7: Explanation ← GenerateXAI(G_hat, L_hat) // feature attribution [8],[21]

8: NetLoad ← L_hat − G_hat

9: a_t ← OptimizeDispatch(NetLoad, X_t, B) // RL policy (DQN/PPO) [7],[9]

10: ExecuteBatteryAction(a_t) // charge / discharge / hold

11: DispatchGrid(a_t) // distribution & delivery

12: B ← UpdateBatteryState(a_t)

13: end loop

Algorithm 1 sums up the closed-loop cycle behind the proposed framework. At every scheduling interval, weather andIoT data are collected, cleaned, and screened for tampering. A topology-aware state estimate and independent generation and demand forecasts are computed, an explainability check is generated, and the reinforcement-learning optimizer decides on the dispatch action that the battery and distribution layers execute before the cycle repeats.

5.7 Computational Complexity and Latency Considerations

Each stage of Algorithm 1 has a different computational cost. Data cleaning and FDI screening are fairly cheap and can run at sub-second latency on commodity hardware. Graph-based, physics-informed state estimation is the most expensive piece of the pipeline. Graph convolution scales with the number of edges in the feeder topology, and evaluating the physics-informed loss term at every node can dominate the per-interval compute budget for a large distribution network [19], [20]. The reinforcement-learning policy, by contrast, is cheap to evaluate at inference time, just a forward pass through a moderately sized network, but expensive to train. That asymmetry supports a practical deployment pattern where policies get retrained periodically offline (seasonally, say, or after major topology changes), while the online control loop simply executes an already-trained policy within the scheduling interval's real-time budget.

6. APPLICATIONS

6.1 Utility-Scale Solar and Wind Farms

At utility scale, improved load forecasting accuracy [2] is expected to benefit day-ahead market bidding by reducing forecast error. Forecast accuracy translates into financial performance in day-ahead electricity markets: a generator that over-forecasts its output has to buy balancing energy to cover the real-time shortfall, exposing it to imbalance penalties, while under-forecasting leaves excess generation that must be curtailed or sold at a less favorable real-time price. LSTM-based forecasts of the kind discussed in Section 4.2 could narrow that exposure at the day-ahead horizon that matters most for bidding. However, this specific link has not been tested in the studies reviewed here and would require market-focused research to confirm.

6.2 Autonomous Microgrids& Hybrid Storage

Reinforcement learning controllers coordinate islanded operation, balancing distributed solar and storage resources [9]. In an islanded, or off-grid, microgrid, there is no external transmission network to absorb imbalances, which raises the stakes of any dispatch decision compared to a grid-connected setting. The DQN and PPO controllers discussed in Section 4.4 are particularly well suited to this setting because they can be trained against a reward function that explicitly penalizes loss-of-load events, not just economic cost [7], [9].

6.3 Smart Cities & EV Charging Infrastructure

Uncoordinated EV charging can overload distribution transformers that were sized decades ago for a much lower-demand era. RL-driven EV charging aggregators address this by dynamically adjusting charging rates based on transformer loading and real-time renewable availability [18]. An aggregator that watches real-time transformer loading and renewable generation can shift flexible charging demand toward periods of high solar or wind output and away from periods of peak transformer stress. That effectively turns EV batteries into a distributed, schedulable load instead of a fixed one [18].

6.4 Industrial Smart Grids & Predictive Maintenance

PMU-assisted predictive maintenance and physics-informed state estimation can help prevent costly downtime in heavy manufacturing facilities [8], [20]. Industrial facilities often run their own internal distribution networks with tight tolerances on voltage quality, since sensitive manufacturing equipment can be damaged or forced offline by even brief voltage excursions. Physics-informed state estimation of the kind discussed in Section 4.3 gives facility operators an early warning of developing voltage or thermal stress before it turns into an actual equipment fault, and federated-learning approaches let predictive-maintenance models be trained across multiple facilities without centralizing sensitive operational data [14].

6.5 Cross-Application Synthesis

Reading Sections 6.1 through 6.4 together, a fairly consistent pattern shows up. Every application domain draws on the same underlying pipeline stages introduced in Section 5, but weights them differently depending on its dominant operational risk. Utility-scale solar and wind farms (Section 6.1) lean most heavily on the forecasting layer, since their biggest exposure is financial, market imbalance penalties, rather than physical instability. Autonomous microgrids (Section 6.2) lean most heavily on the reinforcement-learning dispatch layer, since their biggest exposure is a physical loss-of-load event with no external grid to fall back on. Smart-city EV infrastructure (Section 6.3) leans most heavily on coordinated, multi-agent-style scheduling across many small, flexible loads rather than one large asset. Industrial smart grids (Section 6.4) lean most heavily on the state-estimation and explainability layers, since protecting equipment depends on getting an early, trustworthy warning rather than on economic optimization alone.

This kind of differentiated weighting is a useful design principle for anyone trying to adapt the framework in Section 5 to a specific deployment. Rather than building out every layer of the pipeline with equal effort in every context, the synthesis in Section 6.5 suggests putting initial engineering effort into whichever layer addresses that context's biggest risk first, and treating the remaining layers as a roadmap for later phases. That idea gets made concrete in the phased implementation roadmap in Section 10.6.

7. REAL-WORLD CASE STUDY: MANAGING CALIFORNIA'S DUCK CURVE

To ground the forecasting-to-dispatch logic of the proposed framework in an operating grid rather than only in simulated environments, this section looks at the California Independent System Operator (CAISO), one of the most widely cited real-world examples of AI-relevant, storage-dependent renewable integration.

CAISO coined the term "duck curve" to describe the shape of net load, total demand minus wind and solar generation, on a typical high-solar day. Net load falls sharply through the morning and midday as solar output rises, then ramps steeply upward in the early evening as the sun sets while demand stays high. As utility-scale and behind-the-meter solar capacity has grown, this midday dip has deepened, and on some days has even gone negative, while the evening ramp has gotten steeper. Industry analysts now sometimes describe this pattern as evolving from a "duck" into a "canyon." CAISO has publicly reported curtailing more than two million megawatt-hours of utility-scale wind and solar output annually in recent years just to manage this imbalance, a figure that sits alongside the broader renewable capacity growth trends reported by the IEA [11].

CAISO's response has been a rapid build-out of battery storage, timed specifically to absorb midday solar surplus and discharge it during the evening ramp. Battery capacity on the CAISO grid grew from roughly 500 MW in 2020 to more than 13 GW by early 2025, making it the largest concentration of lithium-ion battery storage on any grid in the world. This build-out closely mirrors the storage-dispatch logic of the framework proposed in Section 5, forecast-driven charging during periods of surplus solar generation and forecast-driven discharging during the evening net-load ramp, though today it is implemented mostly through CAISO's market-based dispatch mechanisms and operator judgment rather than through the fully integrated, explainable, RL-driven pipeline this review proposes.

This case shows both the promise and the current limits of AI-enabled dispatch at grid scale. On one hand, CAISO's experience shows that accurate day-ahead and intraday solar forecasting, combined with large-scale storage, can meaningfully cut curtailment and reshape a genuinely difficult net-load profile. That is the core mechanism the proposed framework's forecasting and RL-dispatch layers are designed to formalize and automate. On the other hand, CAISO's storage dispatch today is coordinated mostly through market bidding and human-supervised operational rules, not through the closed-loop, physics-constrained, explainable pipeline described in Section 5. The fact that multi-million-MWh annual curtailment still happens even with over 13 GW of storage suggests there is still meaningful room for the kind of integrated, forecast-to-dispatch automation this review pulls together from the literature. A fully realized version of the proposed framework, with topology-aware state estimation and RL-based storage dispatch operating in closed loop, looks like a plausible next step beyond CAISO's current market-and-rules-based approach. Testing that claim, though, would require the kind of real-world pilot deployment and evaluation that goes beyond what a literature review can offer.

8. QUANTITATIVE BENEFITS AND ECONOMIC ANALYSIS

8.1 Renewable Utilization Gains

Nagappan et al. [9] report that AI-based approaches for microgrid management consistently improve renewable utilization compared to rule-based baselines across the cases they reviewed, though the magnitude of improvement varies by microgrid configuration. The gain comes mainly from the RL controller's ability to anticipate short-term generation swings and pre-position battery state of charge accordingly, rather than only reacting after an imbalance has already happened.

8.2 Lower Operational Costs

Proactive battery dispatch reduces reliance on expensive peaking power plants [7], [9]. By shifting stored renewable energy into periods that would otherwise need peaking generation, a well-tuned RL dispatch policy cuts both the direct fuel cost of peaking plants and the emissions tied to their operation. These gains are helped along by favorable and improving economics. NREL's bottom-up cost modeling puts 2024 utility-scale, four-hour lithium-ion battery storage costs at roughly $334/kWh, with mid-case projections declining toward the $150 to $240/kWh range by 2035. That directly raises the return on investment from AI-optimized dispatch as the underlying hardware gets cheaper [17].

8.3 Enhanced Grid Reliability & Stability

Automated voltage and frequency control can help stabilize distributed distribution circuits [1], [19], [20]. Combining topology-aware and physics-informed state estimation, as covered in Sections 4.3 and 5.3, gives operators earlier and more physically grounded warning of developing instability than telemetry inspection alone would provide.

8.4 Reduced Operational Carbon Footprint

Minimizing curtailment directly increases clean energy utilization [1], [11]. This suggests a potential carbon benefit: every unit of curtailed renewable generation, of the kind illustrated by CAISO's own curtailment numbers in Section 7, that is instead absorbed by better-timed storage or demand response could displace fossil generation elsewhere on the grid. However, this specific displacement has not been quantified in the studies reviewed here. Fig. 5 visualizes the resulting system gains.

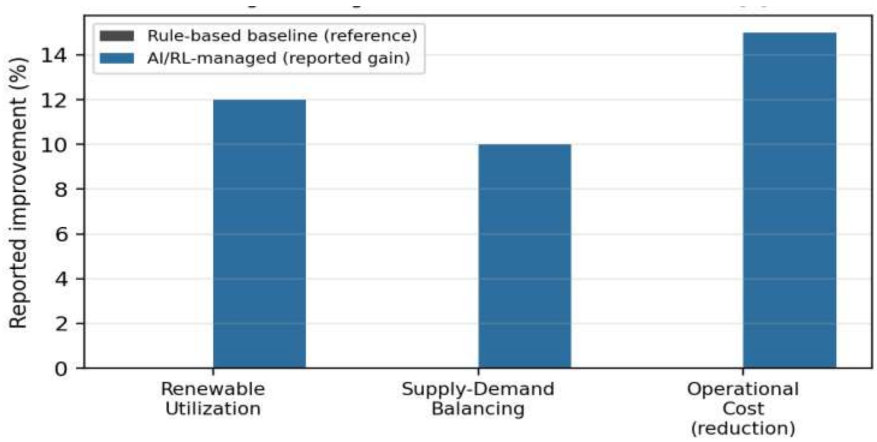

Fig. 5. Illustrative synthesis of reported gains from AI/RL-based microgrid management relative to rule-based control, based on figures reported in Nagappan et al. [9]. Consistent with the limitations noted throughout Table 1, these figures reflect the specific microgrid layout studied and should be treated as indicative rather than universally guaranteed outcomes.

9. CHALLENGES & SECURITY ANALYSIS

9.1 Cybersecurity and False Data Injection (FDI)

Cybersecurity is a major concern, since the same communication infrastructure that enables real-time monitoring and control also widens the attack surface available to malicious actors. Mirzaee, Shojafar, Cruickshank, and Tafazolli's review of smart grid security identifies false data injection, denial-of-service, and eavesdropping as the dominant threat categories, and traces the shift from purely conventional, rule-based defenses toward machine-learning-based detection [15]. Ruan et al. similarly find that deep learning improves intrusion and anomaly detection accuracy compared to signature-based methods, especially for stealthy false-data-injection attacks designed to slip past static thresholds, while also noting that the detection models themselves can be targeted through adversarial or poisoning attacks [16]. Because the proposed framework's state-estimation layer feeds the forecasting ensemble, a successful FDI attack that gets past the initial cleaning step could, in principle, carry corrupted inputs toward the dispatch decision. Ensemble-based deep learning approaches for physical-layer attack localization, like Sakhnini et al.'s, offer one way to mitigate this by flagging measurements that are inconsistent with the network's physical topology instead of relying only on statistical outlier detection [22].

9.2 Model Bias & Transferability

AI models trained on one regional climate or consumer demographic often perform poorly when deployed on a different grid topology without local fine-tuning [4], [8]. A forecasting model trained mostly on data from a temperate, summer-peaking grid might struggle if deployed unmodified on a grid with a winter-peaking heating load or a very different solar resource profile. This limitation shows up in nearly every study summarized in Table 1, and it is a big reason the framework in Section 5 needs periodic retraining as the default assumption, rather than a single static model deployed once and left alone.

9.3 Data Privacy &Infrastructure Constraints

Detailed, meter-level data can reveal private consumer habits, and that is a big part of why privacy-preserving techniques like Federated Learning matter here. Fekri, Grolinger, and Mir show that federated learning with recurrent neural networks applied to distributed smart meter data can roughly match the forecasting accuracy of centrally trained models while keeping household-level data on local devices. The approach still appears vulnerable to non-IID data distributions across nodes, though, which can bias the aggregated global model toward the characteristics of data-richer nodes [14].

9.4 Regulatory, Standardization, and Interoperability Challenges

Beyond the technical challenges above, deploying AI-driven dispatch at utility scale raises governance questions. Alsaigh et al. identify explicit accountability and explainability frameworks as something regulators and grid operators need before they will accept AI-recommended control actions as authoritative [8]. Interoperability makes this harder. The AMI, SCADA, and PMU infrastructure surveyed in Section 3.1 was deployed over decades by different vendors under different standards, and an AI pipeline meant for wide deployment has to reconcile these mismatched data formats and communication protocols rather than assuming one uniform telemetry interface [1], [13].

9.5 Scalability and Computational Cost

The complexity analysis in Section 5.7 pointed out that graph-based and physics-informed state estimation carry computational costs that scale with network size. At the scale of a large regional transmission network, rather than a single feeder or microgrid, this raises open questions about whether the graph-convolution and physics-informed approaches validated on synthetic PMU datasets and smaller test systems will stay tractable in real time without further algorithmic shortcuts [19], [20].

9.6 Workforce and Operational Change Management

A challenge that gets comparatively little attention in the algorithmic literature reviewed in Section 3, but is well documented in the broader smart grid deployment literature, is the human and organizational cost of introducing AI-driven decision support into control rooms staffed by operators trained on conventional SCADA-based procedures [1], [3]. This is not the main focus of this review, but it is worth flagging as a real practical barrier. Even a well-validated forecasting or dispatch recommendation is only useful if operators trust it enough to act on it under time pressure, which is exactly the operational role the explainability layer in Section 5.4 is designed to play. Utilities adopting the framework in Section 5 should budget for operator training and expect an extended period of AI-recommends, human-decides operation before granting any autonomous dispatch authority (Section 10.5), rather than treating workforce adaptation as something to sort out after the technical system is already deployed.

9.7 Standardization and Cross-Vendor Interoperability

Section 9.4 noted that AMI, SCADA, and PMU infrastructure deployed by different vendors over different decades complicates data integration. This gets made worse by the fact that there is currently no widely adopted data-exchange standard specifically for the AI-oriented telemetry, things like model confidence intervals, XAI feature attributions, and GNN topology graphs, that the framework in Section 5 depends on. Existing grid data standards were largely designed before these AI-specific data types even existed. Until such standards mature, utilities piloting the proposed framework should expect to build custom integration layers between the AI pipeline and legacy SCADA/AMI systems, which is itself a nontrivial engineering cost separate from the AI modeling effort discussed in Sections 4 and 5.

10. FUTURE RESEARCH DIRECTIONS

10.1 Digital Twins

Maintaining high-fidelity virtual replicas of physical grid assets could let the AI models and control policies described in Section 5 be tested and refined in simulation before deployment, which would cut operational risk. Jafari, Kavousi-Fard, Chen, and Karimi's review of digital twin technology in smart grids, transportation, and smart cities reports applications spanning real-time condition monitoring, fault diagnosis, and predictive maintenance, and identifies interoperability and real-time synchronization between the physical and virtual asset as the main remaining engineering barriers [12]. Pairing a digital twin with the dispatch pipeline outlined in Section 5 would let a candidate RL policy get tested against simulated extreme-weather scenarios before it is ever allowed to issue a live dispatch action, which directly addresses the sim-to-real gap noted throughout Section 3.6.

10.2 Federated Learning & Edge AI

Training forecasting models across distributed edge nodes can preserve consumer privacy while cutting cloud transmission bandwidth. Extending the federated approach demonstrated by Fekri, Grolinger, and Mir to the state-estimation components of the proposed framework, rather than just the forecasting layer, is a natural next step. It does raise open questions about how to federate a graph-structured model when different edge nodes only see a local subgraph of the full network topology [14].

10.3 Multi-Agent Reinforcement Learning (MARL)

Coordinating multiple autonomous energy aggregators in peer-to-peer trading markets is a natural extension of the single-agent RL formulations reviewed in Sections 3.6 and 4.4 [7], [9]. As distributed energy resources multiply, a single centralized RL optimizer could become both a computational bottleneck and a single point of failure. A MARL formulation, where neighboring microgrids or aggregators each run a local policy while coordinating through a shared market or communication protocol, offers a path toward decentralizing the dispatch problem itself. Shrestha et al.'s review of peer-to-peer energy trading in community microgrids reports increased local renewable self-consumption under these decentralized arrangements, though scalability and transaction-coordination overhead remain open research problems [23].

10.4 Physics-Informed Neural Networks: Next Steps

Embedding physical power-flow constraints directly into model training to guarantee safe grid operation during transients remains an active area of research. This extends beyond the steady-state formulations reviewed in Section 3.4 toward dynamic, transient-stability constraints relevant to fault ride-through and frequency response [20].

10.5 Toward Fully Autonomous Grid Operation

The four directions above point toward a longer-term trajectory: digital twins provide a safe testing ground, federated and edge learning provide privacy-preserving scalability, MARL provides decentralized coordination, and increasingly sophisticated PINNs provide physical guarantees. Together, these four directions move the field from AI as a decision-support tool for human operators toward AI as a trusted, largely autonomous grid-operation layer working under human oversight rather than under human moment-to-moment control. The rules and explainability tools discussed in Sections 3.5 and 9.4 need to keep up with the technical capability, not lag behind it [8], [21].

10.6 Implementation Roadmap

Table 3 in Section 4.6 established that the components of the proposed framework sit at different maturity levels, and Section 6.5 argued that different applications should weight those components differently. Combining these two observations, Table 4 sketches a four-phase implementation roadmap by which a utility or microgrid operator could adopt the framework in Section 5 gradually, rather than attempting a single, all-at-once deployment. This connects the conceptual framework directly to a realistic path toward the real-world validation flagged as missing in Section 12.

TABLE 4 Phased Implementation Roadmap for the Proposed Framework

Phase	Scope	Framework Components	Maturity Basis
1. Forecasting-only pilot	Deploy demand & generation forecasting alongside existing SCADA, in an advisory (non-controlling) role	Weather/IoT ingestion, LSTM/GRU forecasting (5.2–5.3)	Mature components only [2], [5]
2. Add state estimation	Introduce topology-aware, physics-informed state estimation as a parallel monitoring layer	GNN + PINN state estimation (5.3)	Emerging components, still advisory [19], [20]
3. Add explainability& operator interface	Surface feature attributions for all forecasts and state estimates to control-room staff; begin formal operator training (Section 9.6)	XAI layer (5.4)	Governance-ready per [8], [21]
4. Limited	Grant RL-based dispatch	RL optimization	Research-stage;

Phase	Scope	Framework Components	Maturity Basis
autonomous dispatch	authority over a bounded asset set (e.g., a single battery bank), with human override retained	engine (5.5–5.6)	requires digital-twin validation first (10.1)

This roadmap is a starting point for the kind of real-world pilot study that would be needed to test the pipeline described in Section 5, not proof that such a pilot has already happened.

11. DISCUSSION

Looking across the studies reviewed in this paper, no single AI method dominates every smart grid engineering objective. Instead, the methods form a complementary stack where each family makes up for a limitation in another. Classical machine learning is fast and easy to interpret, but it doesn't handle long-range temporal or spatial structure well [2]. Deep sequential models fix the temporal half of that problem, but they still treat the grid as an unstructured pile of time series rather than a physically connected network [2], [5]. Graph neural networks address this spatial limitation by restoring structural awareness, though without physical constraints they can still produce non-physical state estimates under conditions they haven't seen before [19]. Physics-informed learning is what fixes that last part, at the cost of needing accurate network parameters that aren't always available or current [20]. None of this matters much, though, if operators don't trust the output. Explainable AI doesn't improve predictive accuracy at all, but it's what makes the earlier methods usable in a live control room in the first place [8], [21]. Reinforcement learning closes the loop by turning forecasts and state estimates into actual dispatch actions, but it inherits every one of the upstream limitations and adds its own reward-design and sim-to-real headaches on top [7], [9], [18].

This complementary structure is the main reason for combining these methods into a single pipeline rather than treating forecasting, state estimation, explainability, and control as separate problems each solved on its own, which is the pattern identified as the research gap in Section 3.9. The quantitative benefits summarized in Section 8, and the CAISO case study in Section 7, should be read in that light: as evidence that an integrated pipeline can plausibly deliver gains, not as evidence that any single algorithmic component is sufficient on its own.

At the same time, the challenges catalogued in Section 9 don't weigh on the pipeline evenly. Cybersecurity and computational scalability fall hardest on the state-estimation and forecasting layers, since topology-aware and physics-constrained models are the most computationally demanding part of the whole system [15], [16], [19], [20]. Model bias and transferability show up mostly in the forecasting layer, where training data availability differs sharply across regions [4], [5], [6]. Data privacy and regulatory acceptance land more on the data-acquisition and control layers, since consumer-level telemetry and autonomous dispatch actions are what end users and regulators actually see [8], [14]. This uneven spread is a useful way to decide where future engineering and research effort should go first.

11.1 Where the Literature Agrees and Disagrees

Across the sources synthesized in this review, agreement is strongest on two points: that deep learning forecasting methods outperform purely statistical baselines on standard error metrics [2], [5], and that explainability is a prerequisite for, not an optional add-on to, operator and regulatory acceptance of AI-driven control [8], [21]. Disagreement, or more precisely a difference in emphasis, shows up more at the boundary between forecasting-focused and control-focused research. Studies coming out of the power-systems-engineering tradition (for example [19], [20]) tend to prioritize physical consistency and interpretability of state estimates over raw predictive accuracy, while studies coming out of the machine-learning tradition (for example [2], [6]) tend to prioritize accuracy metrics and treat physical plausibility as secondary. Neither position is wrong within its own frame, but the difference in emphasis matters for anyone trying to pick a single method for grid deployment, since the two traditions do not always report results on a metric that allows for direct comparison.

11.2 Emerging Leaders and Areas of Weak Evidence

If forced to say which methods look best positioned for near-term deployment based on the evidence reviewed here, LSTM/GRU forecasting and single-agent RL-based battery dispatch stand out as the most field-proven. That fits their classification in Table 3, which distinguishes between mature single-agent RL applications and research-stage multi-agent RL [2], [5], [7], [9]. Graph neural networks and physics-informed learning for state estimation are methodologically compelling and directly relevant to the framework in Section 5, but they are still supported mainly by evaluations on synthetic PMU datasets and smaller test systems rather than operating utility networks [19], [20]. This is the area of weakest real-world evidence among the technologies reviewed, and the area where the Section 12 caveat about simulation-to-deployment gaps applies most strongly. Multi-agent reinforcement learning and federated learning for grid applications are earlier still, with the reviewed evidence mostly confined to single-study demonstrations rather than a body of independently replicated results [14]. Readers should weight the future-research claims in Sections 10.2 and 10.3 accordingly, as promising directions rather than established practice.

12. LIMITATIONS OF THIS REVIEW

This paper has its own limitations, common to targeted rather than fully systematic literature reviews. First, although the methodology in Section 2 documents the databases, search terms, and inclusion criteria used, this review synthesizes a deliberately narrow set of sources chosen for their relevance to the forecasting-to-dispatch pipeline described in Section 5. It does not claim to be an exhaustive systematic review of the smart grid AI literature, and relevant work outside this specific scope may have been left out. Second, several of the quantitative benefits discussed in Section 8, including the renewable-utilization and cost figures drawn from Nagappan et al. and related studies, are reported within specific microgrid layouts or simulation environments and may not generalize directly to other topologies or scales, as the limitations column of Table 1 makes explicit for nearly every entry [9], [14], [18].

Third, differences in datasets, evaluation metrics, and simulation environments across the reviewed studies make direct quantitative comparison hard. A reported accuracy or utilization gain in one study is not necessarily comparable to a similar-looking figure from another, and this review has tried to flag that non-comparability explicitly rather than lumping figures together across studies. Fourth, as with most technical literature, the reviewed sources may carry some degree of publication bias favoring positive or favorable results, which could cause this review to overstate how effective these methods typically are in the real world compared to unpublished or negative findings. Fifth, the proposed integrated framework in Section 5 is a synthesis grounded in the reviewed literature, not a framework that has actually been implemented and benchmarked end to end as a single system. The computational-complexity discussion in Section 5.7 is analytical rather than measured on physical hardware, and the CAISO case study in Section 7 illustrates the framework's relevance to a real grid without constituting a validation of the framework itself. Finally, AI research moves fast, and new techniques, particularly in graph-based and physics-informed learning, which are still fairly young subfields within power systems, may emerge or mature not long after this review's literature search was conducted. Readers should treat the framework as a structured engineering proposal for future implementation and evaluation, not as a validated deployed system.

13. CONCLUSION

This review looked at classical machine learning, LSTMs and GRUs, graph neural networks, physics-informed learning, explainable AI, and reinforcement learning as tools for turning conventional electricity networks into more resilient, lower-carbon smart grids. Individually, each of these methods is already helping in some corner of grid operations. By unifying forecasting, state estimation, explainability, and dispatch control into one pipeline, the framework proposed here sketches out how grid operators might work toward maximizing renewable utilization while preserving cyber-physical reliability. The literature reviewed points to consistent, if context-specific, gains from pieces of this approach: better forecasting accuracy relative to statistical baselines [2], [5]; more physically consistent state estimation when spatial topology and physical constraints are respected [19], [20]; more operator trust when explainability is built into the pipeline instead of bolted on afterward [8], [21]; and real improvements in renewable utilization and operational cost when reinforcement learning governs dispatch in the studied microgrid contexts [7], [9], [18]. The CAISO case study in Section 7 shows that the forecasting-and-storage logic at the core of this pipeline isn't just theoretical. It already underlies, in a less integrated and less automated form, one of the largest real-world deployments of storage-based renewable integration anywhere in the world.

At the same time, the challenges identified in Section 9 are not side issues to sort out after deployment. Cybersecurity and FDI resilience, model bias and transferability, data privacy, regulatory acceptance, and computational scalability all need to shape the engineering of any AI-driven smart grid system from the very start. Digital twins, federated learning, multi-agent reinforcement learning, and next-generation physics-informed networks, the four directions covered in Section 10, are promising paths toward addressing these constraints while holding onto the benefits described in Section 8. As renewable energy penetration keeps rising worldwide, building explainable, physically-consistent, and adaptive AI into smart grids looks essential for resilient, sustainable, and economically efficient power systems. Not as one single algorithm to deploy, but as the coordinated pipeline this review has tried to pull together from the literature.

14. EVALUATION METRICS

Standard error metrics used across the reviewed forecasting literature include the following [2].

$$\mathrm{MAE} = (1/n)\ \Sigma\ |y_i - \hat{y}_i| \quad (5)$$

$$\mathrm{RMSE} = \sqrt{[(1/n)\ \Sigma\ (y_i - \hat{y}_i)^2]} \quad (6)$$

$$\mathrm{MAPE} = (100\%/n)\ \Sigma\ |(y_i - \hat{y}_i) / y_i| \quad (7)$$

$$R^2 = 1 - [\Sigma(y_i - \hat{y}_i)^2 / \Sigma(y_i - \bar{y})^2] \quad (8)$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, $\bar{y}$ is the mean of the actual values, and n is the number of observations. MAE gives an easily interpretable average error magnitude. RMSE penalizes larger errors more heavily, which is useful for catching models prone to occasional large misses during rapidly changing weather. MAPE expresses error as a percentage, which makes it easier to compare across datasets of different scale. R^2 shows the proportion of variance in the target variable that the model actually explains.

14.1 Worked Example: Interpreting the Metrics Together

To show how these four metrics can diverge in practice, consider a toy five-interval solar forecast with actual outputs y = [40, 42, 5, 44, 45] MW and predictions ŷ = [41, 40, 15, 43, 46] MW, where the third interval represents a sudden cloud-cover event the forecasting model failed to anticipate. The absolute errors are [1, 2, 10, 1, 1], giving MAE = (1+2+10+1+1)/5 = 3.0 MW. Because RMSE squares each error before averaging, that single 10 MW miss dominates the calculation: RMSE = $\sqrt{[(1+4+100+1+1)/5]} = \sqrt{21.4} \approx 4.6$ MW, noticeably larger than MAE even though four of the five intervals were forecast almost perfectly. MAPE, on the other hand, gets distorted in the opposite direction by the low-output interval. The third term contributes |5−15|/5 = 200%, which pulls the average MAPE up sharply even though the absolute error there (10 MW) is the same as it is elsewhere in absolute terms.

This toy example makes a practical point directly relevant to model selection in Sections 4.1 and 4.2. RMSE is the more informative metric when large, infrequent misses (like a missed cloud-cover event) carry disproportionate operational cost, since an unanticipated 10 MW shortfall could force a costly last-minute peaking-plant dispatch that four small 1 MW errors never would. MAPE, on the other hand, becomes unstable and can mislead comparisons whenever actual values get close to zero, which is common for solar output overnight or during heavy cloud cover. A robust evaluation protocol reports MAE, RMSE, and R^2 together rather than relying on MAPE alone, and, where possible, weights errors by their operational consequence, for example by time-of-day electricity price, rather than treating every megawatt of error as equally costly.

REFERENCES

1. M. Kiasari, M. Ghaffari, and H. H. Aly, "A comprehensive review of the current status of smart grid technologies for renewable energies integration and future trends: The role of machine learning and energy storage systems," Energies, vol. 17, no. 16, p. 4128, 2024.
2. B. Biswal, S. Deb, S. Datta, T. S. Ustun, and U. Cali, "Review on smart grid load forecasting for smart energy management using machine learning and deep learning techniques," Energy Reports, vol. 12, pp. 3654–3670, 2024.
3. H. Farhangi, "The path of the smart grid," IEEE Power & Energy Magazine, vol. 8, no. 1, pp. 18–28, 2010.
4. F. G. Bishaw, M. K. Ishak, and T. H. Atyia, "Review Artificial Intelligence Applications in Renewable Energy Systems Integration," Journal of Electrical Systems, vol. 20, no. 3, pp. 566–582, 2024, doi: 10.52783/jes.2983.
5. W. Kong, Z. Y. Dong, Y. Jia, D. J. Hill, Y. Xu, and Y. Zhang, "Short-Term Residential Load Forecasting Based on LSTM Recurrent Neural Network," IEEE Transactions on Smart Grid, vol. 10, no. 1, pp. 841–851, 2017, doi: 10.1109/TSG.2017.2753802.
6. K. Barhmi, C. Heynen, S. Golroodbari, and W. van Sark, "A Review of Solar Forecasting Techniques and the Role of Artificial Intelligence," Solar, vol. 4, no. 1, pp. 99-135, 2024, doi: 10.3390/solar4010005.
7. F. R. Albogamy, S. A. Khan, G. Hafeez, S. Murawwat, S. Khan, S. I. Haider, A. Basit, and K. D. Thoben, "Real-Time Energy Management and Load Scheduling with Renewable Energy Integration in Smart Grid," Sustainability, vol. 14, no. 3, p. 1792, 2022.
8. R. Alsaigh, R. Mehmood, and I. Katib, "AI explainability and governance in smart energy systems: A review," Frontiers in Energy Research, vol. 11, 2023, doi: 10.3389/fenrg.2023.1071291.
9. B. Nagappan, C. Chaudhary, D. K. Doda, and S. Kumar, "Review on Smart Grids and Renewable Integration: An Artificial Intelligence-Powered Perspective," Multidisciplinary Reviews, vol. 6, 2024, doi: 10.31893/multirev.2023SS070.
10. H. AsadiAghajari, T. Niknam, M. Shasadeghi, S. M. Sharifhosseini, M. H. Taabodi, E. Sheybani, G. Javidi, and M. Pourbehzadi, "Analyzing Complexities of Integrating Renewable Energy Sources into Smart Grid: A Comprehensive Review," Applied Energy, vol. 383, p. 125317, 2025.

11. International Energy Agency, "Renewables 2025: Analysis and Forecast to 2030," IEA, Paris, France, 2025.
12. M. Jafari, A. Kavousi-Fard, T. Chen, and M. Karimi, "A review on digital twin technology in smart grid, transportation system and smart city: Challenges and future," IEEE Access, vol. 11, pp. 17471–17484, 2023.
13. S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni, "Smart meters for power grid: Challenges, issues, advantages and status," Renewable and Sustainable Energy Reviews, vol. 15, no. 6, pp. 2736–2742, 2011.
14. M. N. Fekri, K. Grolinger, and S. Mir, "Distributed load forecasting using smart meter data: Federated learning with recurrent neural networks," International Journal of Electrical Power & Energy Systems, vol. 137, p. 107669, 2022.
15. P. H. Mirzaee, M. Shojafar, H. Cruickshank, and R. Tafazolli, "Smart grid security and privacy: From conventional to machine learning issues (threats and countermeasures)," IEEE Access, vol. 10, pp. 52922–52954, 2022.
16. J. Ruan et al., "Deep learning for cybersecurity in smart grids: Review and perspectives," Energy Conversion and Economics, vol. 4, no. 4, pp. 233–251, 2023.
17. W. Cole, V. Ramasamy, and M. Turan, "Cost Projections for Utility-Scale Battery Storage: 2025 Update," National Renewable Energy Laboratory, Golden, CO, USA, NREL/TP-6A40-93281, 2025.
18. H. M. Abdullah, A. Gastli, and L. Ben-Brahim, "Reinforcement learning based EV charging management systems — A review," IEEE Access, vol. 9, pp. 41506–41531, 2021.
19. O. Kundacina, M. Cosovic, and D. Vukobratovic, "State estimation in electric power systems leveraging graph neural networks," in Proc. 2022 17th Int. Conf. Probabilistic Methods Applied to Power Systems (PMAPS), 2022, pp. 1–6, doi: 10.1109/PMAPS53380.2022.9810559.
20. Q.-H. Ngo, B. L. H. Nguyen, T. V. Vu, J. Zhang, and T. Ngo, "Physics-informed graphical neural network for power system state estimation," Applied Energy, vol. 358, p. 122602, 2024.
21. C. Utama, C. Meske, J. Schneider, R. Schlatmann, and C. Ulbrich, "Explainable artificial intelligence for photovoltaic fault detection: A comparison of instruments," Solar Energy, vol. 249, pp. 139–151, 2023.
22. J. Sakhnini, H. Karimipour, A. Dehghantanha, and R. M. Parizi, "Physical layer attack identification and localization in cyber–physical grid: An ensemble deep learning based approach," Physical Communication, vol. 47, p. 101394, 2021.
23. A. Shrestha et al., "Peer-to-peer energy trading in micro/mini-grids for local energy communities: A review and case study of Nepal," IEEE Access, vol. 7, pp. 131911–131928, 2019.

Author Biography

RatibSikder is a 9th-grade student at Berkeley High School with a strong interest in engineering, artificial intelligence, and technology innovation. He is a full-stack web developer who has designed and launched three websites that are currently active on the internet. His interests include machine learning, renewable energy systems, and the application of AI to solve real-world problems. Ratib received the Mayor's Award for Leadership and Extraordinary Efforts in 2026 and earned First Prize at the AABEA Silicon Valley Chapter Science Fair on May 23rd, 2026. He hopes to pursue engineering research and develop technologies that create a positive global impact.

Dr. Abdur R. Sikder is a Faculty in the Department of Computer Science and the Department of Information System at San Francisco State University. He has his Bachelor of Science in Computer Science degree from The University of Waikato, New Zealand. He also has a master's degree in computer science and a Ph.D. in Computer Science from The University of Sydney, Australia. He studied at UC San Diego, Lincoln University, CA (MBA) and he worked for San Francisco Bay University, UC Berkeley, Penn State, MTU, Marquette University, Brac and Dhaka University and life member of AABEA.