

Artificial Intelligence Applications in Power Systems: A Review of Methods, Use Cases, and Future Directions

Yuvraj Dilip Unde, Dhananjay Thorat, Jaya Wagh, Sonali Gosavi, Diba Ansari

Department of Electrical Engineering

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

yuvrajunde.gcoerc@gmail.com

Abstract: Modern electrical networks face growing operational demands as electricity consumption rises and renewable generation, distributed resources, and electric vehicles become more common. Artificial intelligence (AI) can learn useful relationships from operating records and simulated data, supporting tasks such as demand prediction, fault identification, state estimation, stability screening, and energy scheduling. This review organizes reported AI techniques and applications in both large interconnected grids and smaller power systems. It considers rule-based and fuzzy approaches alongside machine learning, deep learning, reinforcement learning, and newer foundation-model methods. The literature suggests that AI may accelerate selected analyses and provide flexible decision support, but dependable performance still relies on representative data, generalization, explainability, cyber protection, and real-world validation. A staged approach that combines AI with engineering models and expert oversight is outlined. This paper synthesizes existing studies and does not claim new experimental results..

Keywords: Artificial intelligence, power systems, machine learning, load forecasting, fault detection, smart grids, stability assessment

I. INTRODUCTION

Electric power networks are changing as renewable generation, energy storage, responsive demand, and digital monitoring are added to conventional plants. These developments make operation more demanding: electrical variables influence one another, network conditions can change quickly, and operators must interpret large data streams under time pressure. Established analytical methods remain essential, although some problems involve heavy computation or specialist judgment. AI can complement these methods by identifying patterns in archived measurements and live data [1], [2].

AI has been investigated in power engineering for many years. Early applications relied on expert knowledge and rule-based systems; later work introduced neural networks, fuzzy reasoning, evolutionary algorithms, and more recently deep-learning and reinforcement-learning methods. Liu and Bose review this evolution from early expert-system prototypes to a wider collection of techniques, while observing that routine industrial adoption continues to face barriers [1].

This review brings together examples from transmission and distribution networks, microgrids, and other small-scale electrical systems. It groups the reported approaches by task, discusses their strengths and limitations, and considers how they might be introduced into practice in a dependable way.



II. REVIEW APPROACH

This narrative review draws on the supplied background paper and key publications covering the history of AI in power engineering [1], critical perspectives on present and future uses [2], applications in small electric power systems [3], AI-assisted analysis of large-grid simulations [4], and distribution-network operation [5]. The material is organized by application area; no new experiment or independently measured performance is claimed. Additional publications are cited for specific methods, evaluation principles, and security topics.

The selected literature addresses prediction, monitoring, control, simulation, renewable-energy integration, and smart-grid management. Because the studies differ in network models, datasets, algorithms, and performance measures, their numerical results cannot be compared directly. The discussion therefore focuses on qualitative findings and practical deployment issues.

III. OVERVIEW OF AI METHODS

A. Expert Systems and Knowledge-Based Methods

Knowledge-based expert systems represent engineering experience as rules and apply an inference process to suggest actions. Their logic can be inspected when rules are clearly stated, making them useful for tasks such as interpreting alarms or supporting restoration. However, building and updating a reliable rule base takes effort, especially as equipment and network conditions evolve [1].

B. Machine Learning and Deep Learning

In supervised learning, a model is trained using examples with known outputs; unsupervised learning instead searches for groups or unusual observations in data without predefined labels. Neural networks and deep-learning architectures can model complex nonlinear and time-varying relationships. Power-system examples include demand and renewable-output prediction, fault recognition, state estimation, and stability screening [1], [3], [5].

C. Reinforcement Learning and Optimization

Reinforcement learning develops a decision policy by letting an agent interact with an environment and receive feedback in the form of rewards. Research has considered its use in energy scheduling, reactive-power management, and coordination of distributed energy resources. Direct experimentation on a live grid may be unsafe, so learning should take place in validated simulators or tightly constrained settings before extensive operational testing [2], [4].

D. Emerging AI Paradigms

Newer AI research includes foundation models, large language models, neuro-symbolic techniques, and systems that combine multiple AI components. They may help engineers search technical knowledge, support documentation, and connect different analytical tasks. Their responses still need independent verification; without validated safeguards, these systems should not be allowed to issue safety-critical grid-control commands on their own [2].

E. Representation, Transfer, and Physics-Informed Learning

Deep-learning architectures contain multiple processing layers that learn useful data representations, reducing reliance on manually designed input features [6]. In power engineering, convolutional and recurrent networks have been explored for waveform recognition, sequential load data, and measurements that vary across time and network location. Such flexibility typically comes with a need for large, representative datasets, careful model tuning, and substantial training resources.

Transfer learning adapts knowledge from one learning task or system to a related target task, which can be valuable when only a small number of labelled examples are available for a new feeder or plant [7]. It is effective only when the source and target settings have meaningful similarities; otherwise, transferred knowledge can reduce performance. Physics-informed neural networks incorporate governing physical relationships into model training [8]. In electrical networks, these constraints may represent power balance, voltage and thermal boundaries, or generator dynamics. Their usefulness depends on whether the embedded physical assumptions accurately describe the real system.



F. Reinforcement-Learning Formulations for Grid Control

A common mathematical description of reinforcement learning is a Markov decision process: the agent observes the current state, selects an action, and receives a reward, aiming to maximize reward over time [9]. For a power grid, state information could include bus voltages, line loading, battery charge, or forecasts. Actions might adjust transformer taps, storage settings, or flexible demand. The reward can combine operating cost and voltage quality while penalizing constraint violations. Earlier studies explored reinforcement learning for stability control [10], and later work reviewed applications in energy management, voltage regulation, market decisions, and demand response [11].

Reward design requires care. If the reward does not reflect actual operating requirements, a controller may optimize its score while choosing an unacceptable action. Safe exploration, explicit treatment of constraints, and realistic simulation are therefore essential. Before a learned policy is considered for operational use, it should be assessed against a conventional controller under comparable conditions.

G. Data Preparation and Feature Engineering

The usefulness of an AI model depends heavily on the quality of its input data. SCADA records, smart-meter readings, phasor measurement units, and weather feeds may differ in sampling interval, units, and error patterns, so their timestamps and formats must be aligned [12]. Common preparation steps include handling outliers and missing values, scaling measurements, and creating features such as previous load values, temperature, calendar variables, and moving averages.

For time-series forecasting, training, validation, and test sets should normally be separated by time so that evaluation reflects genuinely future conditions. Researchers must also avoid data leakage—for example, using a variable that would not be available when the forecast is made. Reporting the preparation procedure makes the study easier to reproduce and audit.

IV. APPLICATIONS IN POWER SYSTEMS**A. Load and Renewable-Energy Forecasting**

Forecasting estimates future demand, renewable output, or market variables from past measurements and factors such as weather, calendar patterns, and current operating conditions. AI models may learn nonlinear and time-dependent behavior that simple rules do not capture well. Better forecasts can inform generation commitment, reserve planning, storage operation, and demand response. Accuracy may fall when the weather, customer behavior, or system configuration differs from the examples used during training [1], [2], [5].

B. Fault Detection and Condition Monitoring

AI models can recognize unusual patterns in voltage, current, frequency, and equipment-condition data, helping with fault identification and maintenance planning. Deep networks can analyze time sequences and signals collected by intelligent electronic devices or phasor measurement units. Real use requires suitable labelled fault examples, careful handling of rare-event imbalance, and testing across fault locations and operating states. AI support must not replace protection equipment designed to meet deterministic safety requirements [3], [5].

C. State Estimation and Stability Assessment

State estimation determines the likely electrical condition of a network using measurements that may be noisy or incomplete. Learning-based estimators can approximate complicated relationships and help identify suspicious measurements. AI has also been studied for voltage and transient-stability screening. Since an incorrect assessment may have serious consequences, the output should be checked using physical constraints, uncertainty information, and fallback procedures [1], [4], [5].

D. Power-Grid Simulation and Decision Support

Analysis of a large grid commonly includes power-flow studies, contingency checks, and transient-stability simulations. Tang et al. describe a framework that combines knowledge representation with patterns learned from simulation data to assist power-flow and stability analysis and transfer knowledge between grids [4]. Such methods



may reduce repetitive calculations and help specialists prioritize cases, but they do not remove the need for validated network models or engineering review.

E. Microgrids and Distributed Energy Resources

Small-scale power systems include isolated and grid-connected systems, microgrids, hybrid installations, and virtual power plants. AI can coordinate local generation, batteries, flexible demand, and renewable resources to support energy management. The result depends on reliable communications, sufficient training examples, and effective coordination between local controllers and the wider network [3].

F. Distribution-System Operation

Distribution networks increasingly accommodate rooftop solar, electric vehicles, storage, and controllable loads. AI has been examined for demand prediction, anomaly recognition, emergency-action support, power-quality monitoring, and operational optimization. Distribution applications must cope with changing network topology, incomplete visibility, and a wide variety of connected devices [5].

G. Demand Response and Energy Management

Demand response changes electricity use in response to prices, grid requirements, or utility requests. AI can estimate how customers might react to incentives and help schedule flexible loads, including water heating, industrial processes, and EV charging. A microgrid energy manager can combine demand and renewable forecasts with battery charge and electricity prices to create a schedule that limits costs and peak demand while using local clean energy effectively. Because these goals can conflict, priorities and constraints should be explicit. Privacy, customer comfort, and communication reliability also matter.

H. Renewable Integration and Storage

Solar and wind output varies with weather and can change rapidly. AI forecasts can help operators schedule reserves and storage, while optimization methods can determine appropriate battery charging and discharging times. A practical controller must account for uncertainty, battery capacity and power limits, degradation, and backup supply needs. Models should be tested across seasons, locations, and unusual operating conditions because performance in one setting may not transfer to another.

I. Electricity Markets and Trading

Participants in electricity markets estimate demand, renewable generation, and price movements when preparing bids and schedules. Machine-learning models can find patterns in historical price and demand records, while optimization can help organize bids and resources. In energy trading, AI is best treated as decision support for forecasting and portfolio analysis, not as a guarantee of future prices. Market rules, network limits, uncertainty, and unexpected events can make historical relationships unreliable; human oversight and regulatory compliance remain necessary.

J. Cybersecurity and Anomaly Detection

Digital communication, remote monitoring, and networked control improve grid visibility but also expose systems to risks such as false-data injection, communication disruption, and unauthorized access. AI-based anomaly detection can flag readings or behavior that depart from normal patterns. An alert alone does not establish that an attack occurred: sensor faults, topology changes, and legitimate unusual events can look similar. Reliable response combines AI results with engineering checks, event logs, and established security procedures. Training data and deployed models also need protection, and both missed events and false alarms should be evaluated.

K. Probabilistic Forecasting and Uncertainty

A point forecast provides one expected value, while a probabilistic forecast represents a range or distribution of possible outcomes and is therefore useful when decisions must account for uncertainty [13]. Quantile regression, ensembles, and distributional neural networks can produce intervals or scenarios for reserve planning, battery scheduling, and risk-aware market participation.

Probabilistic forecasts should be assessed for calibration as well as sharpness. A narrow interval is not useful if it often excludes the actual value. The forecast horizon and update frequency should also be reported, since a model suited to next-hour prediction may not be best for day-ahead planning.



L. Smart-Meter Analytics and Customer Segmentation

Smart meters collect detailed consumption records that can support billing, outage identification, load profiling, and demand-side programs. Their data can be used to forecast consumption at household or aggregate level, group customers with similar patterns, and detect unusual usage [12]. Smart-grid research has also considered deep learning, federated learning, edge intelligence, and distributed computing [14].

Customer groups can help utilities design more suitable tariffs and energy-efficiency programs. However, fine-grained consumption data may reveal household routines. Data aggregation, anonymization, access restrictions, and privacy-preserving learning should be considered when these records are analyzed.

M. Reliability Management and Security Assessment

Reliability management aims to anticipate equipment outages and network conditions that could lead to instability or load shedding. Machine-learning techniques can assist with screening many contingencies and approximating analyses that require substantial computation [15]. In this role, a learned model can prioritize cases that appear critical, after which detailed simulation or engineering analysis is performed. This may reduce workload while retaining a physical check. Missed critical cases and changes in network configuration remain important risks, so conservative thresholds and periodic reassessment are advisable.

N. Electric Vehicles and Flexible Loads

Electric vehicles introduce a flexible load whose timing and size can be uncertain, particularly where feeders and transformers have limited spare capacity. AI can estimate arrival patterns, charging duration, and energy needs from earlier sessions, helping schedule charging to limit peaks and voltage problems. Coordination must respect mobility needs, battery limits, tariffs, and network constraints. Vehicle-to-grid services may add flexibility, but their value depends on battery wear, customer participation, and market rules. Evaluation should include variations in driver behavior rather than relying on one assumed pattern.

V. SYNTHESIS OF THE REVIEWED LITERATURE

The publications offer different but complementary viewpoints. The historical study [1] describes how AI techniques developed and notes the distance between promising research and routine industrial deployment. The critical appraisal [2] emphasizes trustworthy and interpretable methods and considers newer AI approaches. The small-scale power-system review [3] organizes applications around generation, demand, operation, and storage. Tang et al. [4] focus on large-grid simulation data and methods that combine engineering knowledge with learning. Stock et al. [5] review AI for distribution-system operation.

Taken together, these studies suggest that AI is most useful when data volume, nonlinear behavior, or repeated analysis makes an existing task difficult or slow. A model that performs well on one dataset or simulated network, however, is not automatically safe on another. Validation tailored to each application is therefore more reliable than assuming a universal AI solution.

A. Comparison of Common AI Techniques

Different AI approaches suit different tasks. Rule-based systems are comparatively transparent but can be costly to update. Conventional machine learning often works well with structured measurements and may need fewer computing resources than deep networks. Deep learning can process complex time sequences and high-dimensional signals, but it generally requires more data and tuning. Reinforcement learning addresses sequential decisions, although safe constraints must be built into training and evaluation. Transfer learning may reduce data requirements for a related network, but its success depends on how closely the source and target systems match. The method should be chosen according to the engineering goal, data, computing limits, and safety needs.

B. Evaluation Measures

Evaluation measures should reflect the intended engineering task. Forecasts may be compared using mean absolute error, root mean square error, and mean absolute percentage error, noting that percentage error can be misleading when actual values approach zero. For fault recognition, precision, recall, F1-score, confusion matrices, and false-alarm rates



are important, especially when faults are uncommon. For stability screening, missing a dangerous condition can be more costly than issuing a cautious warning. Tests should also report computation time, memory demand, tolerance of missing data, and performance as operating conditions change. Comparisons need a clear baseline and independent test data.

C. Summary of Methods and Typical Tasks

Table I provides a qualitative summary of the methods discussed. It is not intended as a universal ranking: the most suitable approach depends on the application, available data, and operational requirements.

TABLE I: COMPARISON OF AI METHODS FOR POWER-SYSTEM TASKS

Method	Typical tasks	Main strengths	Main limitations
Expert systems / fuzzy logic	Alarm processing, restoration support, uncertain-rule control	Traceable, rule-based reasoning	Knowledge base costly to maintain; rules need tuning
Classical machine learning	Load forecasting, fault classification, anomaly detection	Suits structured data; moderate computing	Needs feature engineering; limited with raw signals
Deep learning	Time-series forecasting, waveform and image-type signal analysis	Learns features automatically	Large data needs; low interpretability
Reinforcement learning	Energy management, voltage and stability control	Handles sequential decisions	Safe exploration; reward design; simulator dependence
Transfer and physics-informed learning	Adaptation to new grids; tasks with scarce data	Lower data need; physical consistency	Depends on source relevance and model accuracy

VI. CHALLENGES AND RESEARCH GAPS

A. Data Quality and Availability

Power-system data may include gaps, sensor errors, imbalanced classes, privacy restrictions, and very few examples of rare but serious events. Better data-sharing arrangements and realistic test cases are needed to assess whether a model will remain reliable in practice.

B. Explainability and Trust

Operators need understandable reasons for a model's recommendation, particularly during abnormal operation. Explainable methods, confidence estimates, and review by human experts may improve trust, but their explanations should be checked for consistency and practical usefulness.

C. Generalization and Transferability

A model trained for one feeder, climate, or operating range may perform poorly in a different setting. Transfer learning and ongoing monitoring may support adaptation, but changes must still pass safety checks. Digital simulation can provide many scenarios for testing before a model is used in the field [4].

D. Cybersecurity and Operational Safety

Manipulated measurements, adversarial inputs, or altered training data can compromise AI systems. Practical deployment should therefore include secure data handling, access controls, continuous model monitoring, independent protection layers, and clear responsibility for approving control actions.

E. Validation and Integration

Future evaluations should document datasets, baseline methods, operating conditions, error measures, computing needs, and failure cases. AI should be combined with established power-system models and assessed through simulation, hardware-in-the-loop testing when appropriate, and carefully staged field trials.



F. Practical Implementation Considerations

A utility should begin with a clearly defined task—for example, short-term load prediction or alarm prioritization—instead of applying a general AI system across the whole network at once. The project should specify who uses the result, what actions it may inform, and what fallback applies when the model is uncertain or unavailable. Data governance should establish ownership, access, retention, and quality checks. Models should first be tested offline and then operated in shadow mode, recording predictions without controlling equipment. A limited operational trial should follow only after satisfactory validation, with monitoring continuing after deployment as sensors, equipment, customer behavior, and topology change.

G. Role of Digital Twins and Simulation

A digital twin or a validated simulation offers a controlled setting for investigating scenarios that may be difficult or unsafe to reproduce on a live network. AI can learn from simulation results, screen many cases, and highlight conditions needing detailed analysis. Tang et al. [4] describe the use of simulation and knowledge discovery for large-grid studies. These benefits depend on simulation fidelity: omitted limits or inaccurate equipment models can create misleading training data. Where measurements are available, simulated behavior should be compared with real operation and discrepancies documented.

H. Evaluating Interpretability

Interpretability has several meanings, so a model should not be called explainable without stating how that claim was assessed. Doshi-Velez and Kim describe application-grounded, human-grounded, and functionally grounded evaluation approaches [16]. For grid operators, application-grounded assessment is especially relevant: the explanation should help them recognize a faulty recommendation or select a better action within the available time. Feature importance, simpler models, and example-based explanations may help in some cases, while safety-critical tasks may require stronger guarantees than a post-hoc explanation alone can provide.

I. Adversarial Manipulation and False Data

Research on false-data injection in state estimation shows that an attacker who understands the network model may alter measurements in ways that evade conventional bad-data checks [17]. AI detectors and estimators can also be misled by crafted inputs or poisoned training records. Possible defenses include redundant measurements, comparison with physical models, authenticated communications, robust training, and monitoring for changes in input data. Because attacks and sensor failures can appear alike, an operator should confirm alerts before automatic corrective actions are taken.

J. Computational Requirements and Real-Time Operation

Grid applications have different response deadlines, from milliseconds for protection-related functions to hours for planning. Large neural models may be accurate but difficult to run on limited substation hardware; very simple models may miss complex patterns. Model compression, edge computing, and hierarchical designs can balance speed and analytical depth—for example, local devices can screen quickly while central servers perform more detailed analysis. Accuracy should be reported alongside latency, memory use, and sensitivity to communication delays.

K. Regulation, Standards, and Human Factors

Utilities must comply with reliability rules, safety requirements, and audit processes, so AI behavior should be documented and traceable. Model and dataset versions, retraining events, and responsibility for decisions should be recorded. Human behavior matters too: frequent false alarms can lead operators to ignore a tool, while apparently strong results can encourage excessive reliance. Training, interface design, and gradual introduction in advisory mode can help develop appropriate trust.

VII. PROPOSED DEPLOYMENT FRAMEWORK

The literature supports a staged deployment process: (1) select a specific engineering task and define its safety limits; (2) assemble and clean representative operating and simulation data; (3) establish a conventional baseline; (4) train and test the AI model on unseen scenarios; (5) add physical checks, uncertainty monitoring, and human approval where



needed; and (6) monitor performance after release and retrain only through a controlled change procedure. This is a framework proposed by the authors from the reviewed work, not a tested prototype.

A forecasting model, for example, should be compared with a statistical or persistence baseline using appropriate error measures. Fault-detection models should be judged by missed faults and false alarms as well as overall accuracy. A stability-screening tool should be tested against validated simulations and retain a conservative fallback when confidence is low.

Consider a distribution utility that needs a day-ahead feeder-load forecast for reserve and storage planning. Historical demand, weather, and calendar records are first cleaned and split into chronological training, validation, and test periods. A seasonal statistical method or persistence forecast provides a baseline. The AI model is then trained and compared with that baseline, including tests on holidays, extreme weather, and unusual demand periods.

Before the forecast affects operations, it can run in shadow mode so that its errors are measured without changing dispatch decisions. Prediction intervals can indicate uncertainty, and the operator can use the baseline when data are missing or uncertainty is high. Following deployment, forecast errors should be monitored and retraining handled through a documented process. This example illustrates the proposed workflow; it does not present measured experimental results.

VIII. FUTURE RESEARCH DIRECTIONS

Further work should prioritize AI that is reliable, explainable, and adaptable across grids. Shared benchmark datasets and transparent test protocols would improve comparison between studies. Hybrid and physics-informed methods could combine electrical laws with learned patterns. Uncertainty-aware models should identify cases outside their experience and refer them for human review. Federated or privacy-preserving learning may allow organizations to benefit from distributed data without centralizing sensitive records. Researchers should also test rare faults, extreme weather, topology changes, and communication failures. Hardware-in-the-loop experiments and staged field studies can help connect simulation results to utility practice.

Federated learning trains a common model across multiple devices or organizations while leaving raw records at their source and exchanging model updates instead [18]. In power systems, this could let utilities, microgrid operators, and customer-side devices contribute to a shared model without transferring detailed consumption data. Remaining challenges include communication overhead, differences between participants' data, and protection against leakage or tampering of model updates.

A further research opportunity is to connect AI continuously with digital twins. As the physical network changes, candidate model updates could be evaluated in a verified simulation before use. Such a process needs explicit rules for retraining, approval, and comparison of each new version with its predecessor. Studies that combine technical measures with feedback from operators would help determine whether AI improves real decisions, rather than only offline scores.

IX. LIMITATIONS OF THIS REVIEW

This article reviews published studies; it does not introduce a new algorithm, laboratory experiment, field trial, or independently collected dataset. The cited works use different applications, assumptions, data, and evaluation procedures, so their results cannot be combined into one numerical performance claim. The discussion of newer AI methods is a broad overview, and each technique requires testing for its intended task. Stronger conclusions about comparative algorithm performance would require a systematic review with predefined search sources, screening rules, and quantitative synthesis.

X. CONCLUSION

AI is increasingly used as a supporting tool for power-system planning, operation, monitoring, and control. The reviewed literature describes applications in forecasting, fault identification, state estimation, stability assessment, simulation studies, microgrid management, and distribution operation. Machine learning and deep learning help



identify patterns, while expert knowledge and reinforcement learning provide complementary capabilities. Their value depends on suitable data, transfer to new conditions, interpretability, cyber resilience, and rigorous validation. Rather than replacing established power-engineering methods, the most dependable approach combines AI with physical models, engineering expertise, and independent safety mechanisms. Future work should emphasize transparent benchmarks, broad testing, and carefully supervised field deployment.

REFERENCES

- [1] C.-C. Liu and A. Bose, "A brief history of application of artificial intelligence to power systems," *Journal of Modern Power Systems and Clean Energy*, vol. 14, no. 1, pp. 1–6, Jan. 2026, doi: 10.35833/MPCE.2025.001004.
- [2] R. J. Bessa, S. Chatzivasileiadis, N. Zhang, C. Kang, and N. Hatziaargyriou, "Current and future applications of artificial intelligence in power systems: A critical appraisal," *Journal of Modern Power Systems and Clean Energy*, vol. 14, no. 1, pp. 23–38, Jan. 2026.
- [3] A. Shahid, F. Plaum, T. Korötko, and A. Rosin, "AI technologies and their applications in small-scale electric power systems," *IEEE Access*, vol. 12, pp. 109985–110012, 2024, doi: 10.1109/ACCESS.2024.3440067.
- [4] Y. Tang, Y. Huang, H. Wang, C. Wang, Q. Guo, and W. Yao, "Framework for artificial intelligence analysis in large-scale power grids based on digital simulation," *CSEE Journal of Power and Energy Systems*, vol. 4, no. 4, pp. 459–468, Dec. 2018, doi: 10.17775/CSEEJPES.2018.01010.
- [5] S. Stock et al., "Applications of artificial intelligence in distribution power system operation," *IEEE Access*, vol. 9, pp. 150112–150131, 2021.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015, doi: 10.1038/nature14539.
- [7] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010, doi: 10.1109/TKDE.2009.191.
- [8] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019, doi: 10.1016/j.jcp.2018.10.045.
- [9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
- [10] D. Ernst, M. Glavic, and L. Wehenkel, "Power systems stability control: Reinforcement learning framework," *IEEE Trans. Power Syst.*, vol. 19, no. 1, pp. 427–435, Feb. 2004.
- [11] D. Cao et al., "Reinforcement learning and its applications in modern power and energy systems: A review," *J. Mod. Power Syst. Clean Energy*, vol. 8, no. 6, pp. 1029–1042, Nov. 2020.
- [12] Y. Wang, Q. Chen, T. Hong, and C. Kang, "Review of smart meter data analytics: Applications, methodologies, and challenges," *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 3125–3148, May 2019, doi: 10.1109/TSG.2018.2818167.
- [13] T. Hong and S. Fan, "Probabilistic electric load forecasting: A tutorial review," *Int. J. Forecast.*, vol. 32, no. 3, pp. 914–938, 2016.
- [14] M. Massaoudi, H. Abu-Rub, S. S. Refaat, I. Chihi, and F. S. Oueslati, "Deep learning in smart grid technology: A review of recent advancements and future prospects," *IEEE Access*, vol. 9, pp. 54558–54578, 2021.
- [15] L. Duchesne, E. Karangelos, and L. Wehenkel, "Recent developments in machine learning for energy systems reliability management," *Proc. IEEE*, vol. 108, no. 9, pp. 1656–1676, Sep. 2020.
- [16] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv:1702.08608*, 2017.
- [17] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Trans. Inf. Syst. Secur.*, vol. 14, no. 1, pp. 13:1–13:33, May 2011.
- [18] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2017, pp. 1273–1282.

