



Paper Type: Original Article

Machine Learning-Driven Optimization of Energy Consumption in Smart Grids: A Comprehensive Review

Pouria Bahramian* 

Department of Computer Engineering, Ayandegan University, Tonekabon, Iran; pouria.bahramian@aihe.ac.ir.

Citation:

Received: 27 January 2024

Revised: 29 March 2024

Accepted: 18 May 2024

Bahramian, P. (2024). Machine learning-driven optimization of energy consumption in smart grids: A comprehensive review. *Metaheuristic algorithms with applications*, 1(2), 143-170.

Abstract

The accelerating complexity of modern power grids, driven by the large-scale integration of renewable energy sources, the proliferation of Distributed Energy Resources (DER), and the increasing volatility of consumer demand patterns, necessitates intelligent and adaptive energy management solutions that surpass the capabilities of conventional optimization approaches. This comprehensive review systematically examines the role of Machine Learning (ML) techniques encompassing supervised learning, unsupervised learning, Reinforcement Learning (RL), and deep learning in optimizing energy consumption across smart grid infrastructures. We provide an in-depth analysis of ML applications spanning Demand-Side Management (DSM), short-term and Long-Term Load Forecasting (LTLF), renewable energy integration, fault detection and diagnostics, energy storage optimization, and real-time adaptive grid control. This paper reviews and synthesizes findings from over 100 studies published between 2018 and 2025, categorizing ML approaches by their application domain, algorithmic paradigm, and deployment context. We evaluate their performance using standardized metrics including Mean Absolute Percentage Error (MAPE), root mean squared error, cost reduction ratios, and Peak-to-Average Ratio (PAR) improvements. Critical research gaps are identified, including challenges related to training data scarcity and quality, model interpretability in safety-critical infrastructure, scalability from pilot projects to grid-scale deployment, computational constraints for real-time inference, and privacy concerns arising from granular smart meter data. Furthermore, we provide a forward-looking analysis of emerging trends that are poised to reshape the field, including Federated Learning (FL) for privacy-preserving collaborative model training, transfer learning for cross-domain generalization, edge-AI deployment for low-latency on-device intelligence, foundation models for energy time-series, and explainable AI frameworks for regulatory compliance. This review serves as a comprehensive reference for researchers, power systems engineers, and policymakers seeking to understand the current state-of-the-art and chart future research directions at the intersection of ML and smart grid energy optimization.

Keywords: Machine learning, Smart grid, Energy optimization, Load forecasting, Demand-side management, Deep learning, Reinforcement learning, Renewable energy integration, Energy efficiency, Federated learning.

1 | Introduction

The global energy landscape is undergoing a profound transformation driven by the convergence of climate change imperatives, rapid urbanization, population growth, and the accelerating deployment of renewable energy technologies. According to the International Energy Agency (IEA), global electricity demand is

 Corresponding Author: Pouria.bahramian@aihe.ac.ir

 <https://doi.org/10.48313/maa.v1i2.73>



Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0>).

projected to increase by more than 30% by 2030 relative to 2020 levels, placing unprecedented strain on aging grid infrastructures that were designed for unidirectional power flow from centralized generation facilities [1]. Simultaneously, the urgency to decarbonize the power sector, which accounts for approximately 40% of global carbon dioxide emissions, has catalyzed the integration of Variable Renewable Energy Sources (VRES) such as solar photovoltaics and wind turbines, whose intermittent and stochastic generation profiles introduce significant operational challenges for grid operators tasked with maintaining supply-demand balance at all times [2].

The traditional power grid, conceptualized in the early twentieth century, was architected as a hierarchical, centralized system wherein large-scale generation plants transmitted electricity through high-voltage transmission networks and stepped it down through distribution subsystems for end-consumer delivery. This legacy architecture, while robust for its era, lacks the flexibility, observability, and responsiveness required to accommodate the decentralized, bidirectional, and highly dynamic energy ecosystem of the twenty-first century [3]. The emergence of the smart grid paradigm represents a fundamental reimagining of this infrastructure, integrating advanced Information and Communication Technologies (ICT) with the physical power delivery system to create a cyber-physical energy network capable of self-monitoring, self-healing, and adaptive optimization [4].

A modern smart grid is characterized by several defining architectural components. Advanced Metering Infrastructure (AMI) provides high-resolution, bidirectional communication between utilities and consumers, enabling granular monitoring of consumption patterns at intervals as fine as one second. Supervisory Control and Data Acquisition (SCADA) systems facilitate centralized monitoring and control of generation and transmission assets. Phasor Measurement Units (PMUs) deliver synchronized, high-frequency measurements of voltage and current phasors across the wide-area network, enabling real-time state estimation and situational awareness. An expanding ecosystem of Internet of Things (IoT) sensors, deployed across generation sites, distribution feeders, substations, and consumer premises, generates massive volumes of heterogeneous, time-stamped data streams that collectively constitute the informational backbone of smart grid intelligence [5].

However, the realization of truly intelligent smart grid operations faces formidable challenges that span multiple dimensions. The intermittency and non-dispatchability of renewable energy sources introduce volatility in generation schedules, necessitating rapid ramping of conventional generators and sophisticated forecasting capabilities. Demand-side volatility, driven by the electrification of transportation, heating, and industrial processes, creates increasingly unpredictable load profiles. Grid stability concerns, including frequency deviations, voltage fluctuations, and cascading failures, are amplified by the reduction in system inertia as synchronous generators are displaced by inverter-based resources. Cybersecurity vulnerabilities in the interconnected cyber-physical infrastructure present risks of coordinated attacks capable of disrupting critical energy services. Furthermore, the sheer volume, velocity, and variety of data generated by smart grid sensors overwhelm traditional rule-based and analytical approaches to grid management [6].

Classical optimization methods, including linear programming, mixed-integer programming, convex optimization, and dynamic programming, have historically formed the mathematical foundation for power systems optimization problems such as Optimal Power Flow (OPF), unit commitment, and economic dispatch. While these methods offer provable optimality guarantees under well-defined assumptions, they suffer from critical limitations in the smart grid context. Specifically, they require explicit mathematical formulations of objective functions and constraints that may not adequately capture the nonlinear, non-convex, stochastic, and high-dimensional nature of real-world energy systems. Their computational complexity often scales exponentially with problem size, rendering them intractable for real-time applications involving millions of decision variables. Moreover, they cannot learn from historical data and adapt to evolving system conditions without manual model recalibration [7].

Machine Learning (ML), by contrast, offers a fundamentally different paradigm: the ability to automatically extract complex patterns, relationships, and decision policies from vast quantities of historical and real-time

data without requiring explicit programming of the underlying physics. ML models can capture nonlinear dynamics, handle high-dimensional input spaces, adapt to concept drift, and provide predictions and control decisions at computational speeds compatible with real-time grid operation. The past decade has witnessed an explosion of research applying ML techniques to virtually every aspect of smart grid optimization, from load forecasting and renewable generation prediction to demand response optimization, fault detection, and autonomous grid control [8].

Despite this growing body of literature, several critical gaps motivate the present review. First, existing surveys tend to focus on specific application domains (e.g., load forecasting or demand response) without providing a unified perspective across the full spectrum of ML-driven smart grid optimization. Second, the rapid emergence of new ML paradigms including Federated Learning (FL), graph neural networks, transformer architectures, and Large Language Models (LLMs) necessitates an updated synthesis that captures the latest methodological advances and their applicability to energy systems. Third, practical deployment challenges related to data quality, model interpretability, scalability, privacy, and cybersecurity are often underexplored in technically focused reviews.

This review addresses the following research questions:

- I. What are the primary ML paradigms and architectures that have been applied to smart grid energy optimization, and how do they compare in terms of predictive accuracy, computational efficiency, and practical deployability?
- II. How have ML techniques been applied across the major application domains of load forecasting, Demand-Side Management (DSM), renewable energy integration, grid stability, and real-time control?
- III. What are the critical barriers to the large-scale adoption of ML-driven optimization in production smart grid environments?
- IV. What emerging ML trends and research directions hold the greatest promise for advancing the field?

The remainder of this paper is organized as follows. Section 2 provides essential background on smart grid architecture, ML paradigms, and evaluation methodologies. Section 3 reviews ML-driven load forecasting and demand prediction. Section 4 examines DSM and energy efficiency applications. Section 5 discusses renewable energy integration and grid stability. Section 6 focuses on Reinforcement Learning (RL) for real-time grid control. Section 7 addresses privacy, security, and FL. Section 8 explores emerging trends and future directions. Section 9 concludes the review with a summary of key findings and recommendations for future research.

2 | Background and Preliminaries

2.1 | Smart Grid Architecture and Data Infrastructure

The smart grid can be conceptualized as a layered cyber-physical system comprising four principal operational layers: generation, transmission, distribution, and consumption. Each layer generates distinct categories of data, operates under different temporal and spatial scales, and presents unique optimization challenges amenable to ML-based solutions.

The generation layer encompasses large-scale centralized power plants (thermal, nuclear, and hydroelectric), utility-scale renewable energy installations (solar farms and wind parks), and increasingly distributed generation resources including rooftop photovoltaic systems, small-scale wind turbines, and combined heat and power units. Data generated at this layer includes generator output power, fuel consumption rates, turbine rotational speeds, inverter status signals, irradiance and wind speed measurements at renewable sites, and real-time market price signals. The primary optimization objectives at the generation layer include economic dispatch, unit commitment, and OPF, all of which determine the most cost-effective and reliable combination of generation resources to meet anticipated demand [9].

The transmission layer transports bulk electricity from generation sites to load centers through high-voltage (110 kV to 765 kV) transmission lines, substations, and interconnection points. PMUs deployed at critical transmission nodes provide synchronized measurements at rates of 30 to 120 samples per second, enabling wide-area monitoring and dynamic security assessment. SCADA systems at substations collect voltage, current, power flow, transformer tap position, and circuit breaker status data at intervals of two to ten seconds. Transmission-level optimization concerns include congestion management, voltage stability, transient stability, and the management of inter-area power flows across interconnected balancing areas [10].

The distribution layer delivers electricity from transmission substations to end consumers through medium-voltage (4 kV to 35 kV) and low-voltage (120 V to 480 V) feeders. The distribution network is the most geographically extensive and topologically complex component of the power system, comprising millions of line segments, transformers, switches, capacitor banks, and voltage regulators. Distribution automation systems, equipped with Intelligent Electronic Devices (IEDs) and Remote Terminal Units (RTUs), provide increasing observability of network conditions. The integration of Distributed Energy Resources (DERs), including rooftop solar, Battery Energy Storage Systems (BESS), and Electric Vehicle (EV) chargers, at the distribution level transforms traditionally passive loads into active prosumers capable of both consuming and injecting power, significantly complicating distribution network management [11].

The consumption layer interfaces directly with end-users across residential, commercial, and industrial sectors. Smart meters, the cornerstone of AMI, record electricity consumption at intervals typically ranging from 15 minutes to one hour, though some advanced deployments support sub-minute granularity. Home Area Networks (HANs) connect smart meters with in-home displays, smart thermostats, smart appliances, and Home Energy Management Systems (HEMS), enabling automated load control and demand response participation. Building Energy Management Systems (BEMS) in commercial and institutional buildings monitor and control Heating, Ventilation, Air Conditioning (HVAC), lighting, and plug loads. Industrial energy management involves process-level monitoring and optimization of manufacturing equipment, compressed air systems, and electric arc furnaces [12].

The data generated across these four layers exhibits several characteristics that make it particularly amenable to and challenging for ML analysis. First, smart grid data is inherently high-dimensional, with each consumer's consumption profile influenced by dozens of variables including weather conditions, time of day, day of week, seasonal patterns, building characteristics, occupancy, and electricity prices. Second, the data exhibits strong temporal dependencies across multiple time scales, from sub-second transients to daily, weekly, and seasonal cycles. Third, data quality issues including missing values, measurement noise, sensor drift, communication failures, and outliers are pervasive in real-world deployments. Fourth, the heterogeneity of data sources spanning numerical time series, categorical event logs, text-based maintenance records, and spatial network topology data requires sophisticated data fusion and preprocessing pipelines [13].

Communication infrastructure is the enabling fabric that connects the physical and cyber layers of the smart grid. Wide Area Networks (WANs) interconnect control centers, substations, and generation facilities using fiber optic, microwave, and cellular communication technologies. Field Area Networks (FANs) connect distribution automation devices and AMI data concentrators. HANs employ short-range wireless protocols such as Zigbee, Z-Wave, Wi-Fi, and Bluetooth to interconnect smart devices within consumer premises. The latency, bandwidth, reliability, and security characteristics of these communication networks directly influence the feasibility of real-time ML-based control applications and the design of distributed learning architectures [14].

2.2 | Machine Learning Paradigms for Energy Systems

ML paradigms applicable to smart grid optimization can be broadly categorized into supervised learning, unsupervised learning, deep learning, RL, and hybrid approaches. Each paradigm offers distinct advantages and limitations depending on the nature of the available data, the optimization objective, the required prediction horizon, and the real-time computational constraints.

Supervised learning encompasses algorithms that learn a mapping from input features to output labels using a training dataset of labeled examples. In the regression setting, supervised learning is extensively applied to load forecasting, where historical consumption data paired with exogenous features (temperature, humidity, solar irradiance, day type, holiday indicators) are used to train models that predict future energy demand. Prominent regression algorithms in smart grid applications include Support Vector Regression (SVR), which constructs an optimal hyperplane in a high-dimensional feature space to minimize prediction error within an epsilon-insensitive tube [15]; Random Forests (RF), which aggregate predictions from an ensemble of decorrelated decision trees to achieve robust, low-variance estimates [16]; Gradient Boosting Machines (GBM), including XGBoost and LightGBM, which sequentially train weak learners to correct the residual errors of preceding models [17]; and linear regression variants including Ridge, Lasso, and Elastic Net regularization for high-dimensional feature spaces with collinearity. In the classification setting, supervised learning is applied to fault detection and diagnosis, power quality event classification, and islanding detection. Support Vector Machines (SVM) with various kernel functions (RBF, polynomial) provide effective nonlinear classification boundaries. Decision tree ensembles, k-Nearest Neighbors (k-NN), and logistic regression serve as baseline classifiers frequently benchmarked against more complex architectures [18].

Unsupervised learning algorithms identify patterns, groupings, and anomalies in unlabeled data without explicit supervision. Clustering algorithms are widely used for consumer segmentation, wherein utility customers are grouped based on the similarity of their consumption profiles to enable targeted demand response programs, differentiated pricing schemes, and network planning. K-means clustering partitions consumers into k non-overlapping groups by minimizing within-cluster variance; however, it requires a priori specification of the number of clusters and assumes spherical cluster geometries. Density-Based Spatial Clustering of Applications with Noise (DBSCAN) identifies clusters of arbitrary shape based on local density and naturally handles noise points, making it suitable for detecting atypical consumption behaviors. Hierarchical clustering produces a dendrogram of nested clusters at varying granularity levels, enabling multi-resolution consumer analysis. Gaussian Mixture Models (GMMs) provide a probabilistic framework for soft cluster assignments. Additionally, unsupervised anomaly detection techniques, including isolation forests, one-class SVMs, and autoencoders, are employed for detecting non-technical losses (energy theft), equipment malfunctions, and cyberattacks in smart meter data streams [19].

Deep learning architectures have demonstrated transformative performance gains across multiple smart grid applications by automatically learning hierarchical feature representations from raw data without manual feature engineering. Convolutional Neural Networks (CNNs), originally designed for image recognition, have been adapted for spatial feature extraction from grid topology data, multi-sensor arrays, and frequency-domain representations of electrical signals. One-dimensional CNNs applied to time-series consumption data capture local temporal patterns and periodic structures through learnable convolutional filters [20]. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), are the predominant architectures for temporal sequence modeling in energy systems. LSTMs address the vanishing gradient problem inherent in vanilla RNNs through gated memory cells that selectively retain, forget, and output information across long time horizons, making them well-suited for capturing the multi-scale temporal dependencies characteristic of energy consumption data [21]. The Transformer architecture, introduced by Ashish [22], has emerged as a powerful alternative to recurrent models for sequence-to-sequence learning. The self-attention mechanism enables Transformers to capture long-range dependencies without sequential processing, offering superior parallelization and scalability. Temporal Fusion Transformers (TFT), specifically designed for multi-horizon time-series forecasting, incorporate variable selection networks, gated residual connections, and interpretable multi-head attention to provide both accurate predictions and insights into feature importance [23]. Autoencoders, including Variational Autoencoders (VAEs) and denoising autoencoders, serve dual roles in dimensionality reduction and generative modeling, enabling compressed representations of high-dimensional energy data and synthetic data generation for augmenting scarce training datasets [24].

RL provides a framework for sequential decision-making in which an agent learns an optimal policy by interacting with an environment and receiving scalar reward signals. RL is uniquely suited for real-time control and optimization problems in smart grids where decisions must be made sequentially under uncertainty and the consequences of actions unfold over extended time horizons. Tabular Q-learning, applicable to small, discrete state-action spaces, maintains a value function table that converges to the optimal policy under sufficient exploration. For the high-dimensional, continuous state-action spaces typical of power systems, Deep Reinforcement Learning (DRL) methods are essential. Deep Q-Networks (DQN) approximate the Q-function using neural networks, enabling generalization across unseen states [25]. Policy gradient methods, including reinforce and actor-critic architectures (A2C, A3C), directly parameterize and optimize the policy function, accommodating continuous action spaces required for voltage setpoint adjustment, generator dispatch, and battery charge/discharge control. Proximal Policy Optimization (PPO) introduces a clipped surrogate objective that constrains policy updates to a trust region, providing stable and sample-efficient training [26]. Soft Actor-Critic (SAC) maximizes a maximum-entropy objective, encouraging exploration and robustness to environmental perturbations [27].

Hybrid and ensemble approaches combine multiple ML paradigms to leverage complementary strengths and mitigate individual weaknesses. CNN-LSTM hybrids, for example, employ convolutional layers to extract local spatial or frequency-domain features that are subsequently processed by LSTM layers for temporal sequence modeling, achieving superior load forecasting accuracy compared to either architecture in isolation [28]. Stacking ensembles trains a meta-learner on the outputs of diverse base models (e.g., SVR, random forest, LSTM) to produce optimally weighted predictions. Physics-Informed Neural Networks (PINNs) embed physical laws such as Kirchhoff's laws, power balance equations, or thermodynamic constraints into the neural network loss function, improving generalization in data-scarce regimes and ensuring physically plausible predictions [29].

2.3 | Evaluation Metrics and Benchmarking

The rigorous evaluation and fair comparison of ML models for smart grid applications require standardized metrics and benchmark datasets. Forecasting performance is typically assessed using several complementary metrics. Mean Absolute Error (MAE) quantifies the average magnitude of prediction errors in the original units of measurement, providing an intuitive and easily interpretable accuracy measure. Root Mean Squared Error (RMSE) penalizes large prediction errors more heavily than MAE due to the squaring operation, making it particularly relevant for grid operations where extreme forecast errors can trigger costly corrective actions. Mean Absolute Percentage Error (MAPE) normalizes errors relative to actual values, enabling comparison across load profiles of different magnitudes; however, MAPE is undefined when actual values are zero and can be biased toward under-prediction. The coefficient of determination (R^2) measures the proportion of variance in the target variable explained by the model, with values approaching unity indicating near-perfect fit [30].

Classification tasks, such as fault detection, power quality event identification, and islanding detection, are evaluated using accuracy, precision, recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC). In smart grid applications where the costs of false negatives (missed faults) and false positives (unnecessary alarms) are asymmetric, precision-recall tradeoffs and F1-scores are more informative than overall accuracy, particularly in class-imbalanced scenarios where fault events constitute a small fraction of total observations [31].

Optimization performance metrics assess the tangible impact of ML-driven solutions on grid operations. Cost reduction percentage measures the savings achieved relative to baseline (non-ML) control strategies. Peak-to-Average Ratio (PAR) quantifies the effectiveness of demand response and load shifting programs in flattening the demand curve. Carbon emission reduction captures the environmental benefits of ML-optimized renewable integration and dispatch. Energy Not Served (ENS) and Loss of Load Probability (LOLP) measure reliability improvements. Computational metrics, including training time, inference latency, and model size,

are critical for evaluating the real-time deployability of ML solutions on resource-constrained edge devices [32].

Several publicly available benchmark datasets have become standard testbeds for smart grid ML research. The Reference Energy Disaggregation Dataset (REDD) provides whole-house and circuit-level electricity measurements from six residences at one-second granularity, widely used for Non-Intrusive Load Monitoring (NILM) algorithm evaluation [33]. The UK Domestic Appliance-Level Electricity (UK-DALE) dataset extends REDD with five UK houses monitored over multiple years [34]. The Pecan Street Dataport contains minute-level data from over 1,000 homes in Austin, Texas, including solar generation and EV charging. The Commission for Energy Regulation (CER) Ireland dataset provides half-hourly smart meter readings from approximately 5,000 residential customers over 18 months, commonly used for consumer segmentation and demand response studies. The Australian Energy Market Operator (AEMO) dataset offers five-minute resolution system-level demand data for the Australian national electricity market [35].

Table 1. Summary of common ML techniques and their primary smart grid applications.

ML Technique	Category	Primary Applications	Key Strengths	Key Limitations
SVR/RF	Supervised	Load forecasting, fault classification	Robust, interpretable, low data requirements	Limited capacity for complex temporal patterns
XGBoost / LightGBM	Supervised	Load forecasting, feature ranking	High accuracy, handles missing data, fast training	Requires careful hyperparameter tuning
LSTM/GRU	Deep learning	Short-term load forecasting, price prediction	Captures long-term temporal dependencies	Sequential processing, slow training
CNN-LSTM Hybrid	Deep learning	Spatio-temporal forecasting, NILM	Combines spatial and temporal feature extraction	Complex architecture, large data needs
Transformer / TFT	Deep learning	Multi-horizon forecasting	Parallel processing, attention interpretability	High computational cost, data-hungry
K-means/ DBSCAN	Unsupervised	Consumer segmentation, anomaly detection	Scalable, no labels required	Sensitive to parameters, no temporal modeling
Autoencoder / VAE	Deep learning	Anomaly detection, data compression	Unsupervised, learns latent representations	Reconstruction quality depends on architecture
DQN/PPO/SAC	RL	Real-time control, energy trading, scheduling	Adaptive, sequential decision-making	Sample inefficiency, safety constraints hard to enforce
MARL	RL	DER coordination, P2P trading	Distributed, scalable to large systems	Non-stationarity, communication overhead
FL	Distributed ML	Privacy-preserving forecasting	Data stays on-device, regulatory compliance	Non-IID data, communication bottleneck

3 | Machine Learning-Driven Load Forecasting and Demand Prediction

3.1 | Short-Term Load Forecasting

Short-term load forecasting, spanning prediction horizons from one hour ahead to one week ahead, constitutes the most extensively studied application of ML in smart grids. Accurate STLTF is indispensable for a multitude of operational decisions including real-time generation dispatch, spinning reserve allocation, energy market bidding, and transmission congestion management. A one-percentage-point improvement in MAPE for a large utility can translate to millions of dollars in annual savings through reduced ancillary service procurement and more efficient generator scheduling [36].

Traditional statistical methods for STLTF, including Autoregressive Integrated Moving Average (ARIMA), Seasonal Autoregressive Integrated Moving Average (SARIMA), and exponential smoothing state space models, have established well-understood baselines. ARIMA models capture linear temporal dependencies through autoregressive and moving average components applied to stationary time series, while SARIMA extends this framework to accommodate periodic seasonal patterns. While computationally efficient and statistically interpretable, these methods fundamentally assume linear relationships between past and future observations and struggle to incorporate high-dimensional exogenous variables, limiting their accuracy when

load dynamics are driven by complex, nonlinear interactions among weather, occupancy, pricing, and behavioral factors [37].

The application of LSTM networks to STLF has demonstrated consistent and substantial improvements over statistical baselines. Kong et al. [38] proposed a deep LSTM architecture for residential STLF that achieved a MAPE of 5.2% on the Pecan Street dataset, outperforming ARIMA (MAPE 12.7%), SVR (MAPE 8.4%), and shallow feedforward neural networks (MAPE 7.1%). The gated memory cell architecture of LSTMs enables effective learning across multiple temporal scales, from hourly fluctuations driven by occupancy transitions to weekly patterns reflecting workday-weekend distinctions. Bidirectional LSTMs (Bi-LSTMs), which process input sequences in both forward and reverse temporal directions, have been shown to further improve forecasting accuracy by capturing both causal and anti-causal dependencies in the data [36].

CNN-LSTM hybrid architectures have emerged as a particularly effective paradigm for STLF by synergistically combining the spatial feature extraction capabilities of convolutional layers with the temporal sequence modeling strengths of recurrent layers. In a representative study, Kim and Cho [39] designed a two-stage architecture wherein one-dimensional convolutional layers first extract local patterns and reduce dimensionality from multi-variate input features (including load, temperature, humidity, wind speed, and solar irradiance), and stacked LSTM layers subsequently model the temporal evolution of these extracted features. Evaluated on the Korean Electric Power Corporation (KEPCO) dataset, this hybrid architecture achieved a MAPE of 2.13% for day-ahead forecasting, representing a 23% relative improvement over standalone LSTM models.

Attention mechanisms, originally proposed for neural machine translation, have been adapted for load forecasting to enable models to selectively focus on the most informative time steps and features when generating predictions. The TFT, proposed by Lim et al. [23], represents a state-of-the-art architecture specifically designed for multi-horizon time-series forecasting. The TFT incorporates several specialized components: variable selection networks that automatically identify the most relevant input features at each time step, gated residual networks that enable skip connections and controlled information flow, and interpretable multi-head attention that reveals which historical time steps most influence future predictions. Applied to electricity demand forecasting, the TFT achieved improvements of 7-12% in normalized quantile loss compared to competitive baselines including DeepAR, N-BEATS, and standard Transformer models. Critically, the TFT's built-in interpretability mechanisms provide grid operators with transparent explanations of forecast drivers, addressing a significant barrier to ML adoption in safety-critical infrastructure [40].

The incorporation of exogenous features, particularly meteorological variables, is critical for accurate STLF, given the strong influence of weather conditions on temperature-sensitive loads such as HVAC systems, which can account for 40-60% of building energy consumption. Advanced feature engineering approaches include the construction of weather interaction terms (e.g., heating degree days, cooling degree days, wind chill indices), the encoding of calendar effects using cyclical features (sine and cosine transformations of hour, day-of-week, and month), the inclusion of economic indicators (electricity price, GDP growth), and the integration of social event calendars. Recent work has explored the use of Numerical Weather Prediction (NWP) ensemble forecasts as inputs to ML models, where the ensemble spread provides implicit uncertainty information that can improve both point forecast accuracy and prediction interval reliability [41].

3.2 | Medium-Term and Long-Term Load Forecasting

Medium-Term Load Forecasting (MTLF), covering horizons from one month to one year, and Long-Term Load Forecasting (LTLF), extending from one year to several decades, serve fundamentally different operational purposes than STLF. MTLF supports fuel procurement planning, maintenance scheduling, hydrothermal coordination, and medium-term financial planning. LTLF informs capacity expansion planning, infrastructure investment decisions, regulatory filings, and energy policy formulation. The accuracy requirements and methodological approaches for MTLF and LTLF differ substantially from STLF due to the

reduced relevance of short-term weather fluctuations and the increasing importance of macroeconomic, demographic, and technological trend factors [42].

Transfer learning has emerged as a valuable technique for addressing data scarcity challenges in load forecasting, particularly for newly constructed buildings, emerging markets, and developing regions where historical consumption data is limited or unavailable. The fundamental premise is that models trained on data-rich source domains can transfer learned representations to data-scarce target domains, thereby accelerating learning and improving generalization. Ribeiro et al. [43] demonstrated that a CNN-LSTM model pre-trained on three years of aggregated load data from a large European utility could be fine-tuned with as few as two weeks of data from a new building to achieve MAPE values within 2% of models trained on the full target dataset. Domain adaptation techniques, including Maximum Mean Discrepancy (MMD) minimization and adversarial domain adaptation, have been applied to align the feature distributions between source and target domains, further improving transfer effectiveness across geographically and climatically distinct regions [44].

Ensemble methods that combine predictions from multiple heterogeneous models consistently outperform individual models for MTLF by reducing variance and capturing diverse aspects of the underlying data-generating process. Gradient boosting ensembles, particularly XGBoost and LightGBM, have proven exceptionally effective for MTLF due to their native handling of missing values, built-in regularization, and computational efficiency. A comprehensive comparative study by Haben et al. [45] evaluated 42 forecasting methods on the CER Ireland dataset and found that gradient boosted tree ensembles achieved the lowest MAE for monthly forecasting horizons, while LSTM-based models excelled at capturing sudden structural changes in demand patterns.

A critical challenge for long-horizon forecasting is concept drift, the gradual or sudden change in the statistical properties of the target variable over time. In the energy domain, concept drift can arise from energy efficiency improvements, building retrofits, changes in occupancy patterns, deployment of distributed generation, adoption of EVs, and macroeconomic shifts. The COVID-19 pandemic of 2020-2021 provided a dramatic illustration of this challenge, with national electricity demand in many countries dropping by 10-30% within weeks as lockdown measures curtailed commercial and industrial activity while increasing residential consumption. Standard ML models trained on pre-pandemic data exhibited MAPE increases of 50-200% during the initial lockdown period. Adaptive approaches, including online learning with exponential forgetting, ensemble methods with dynamic weight adjustment, and change-point detection algorithms coupled with model retraining triggers, have been proposed to address concept drift [46].

3.3 | Probabilistic and Uncertainty-Aware Forecasting

Point forecasts, while widely used, provide an incomplete picture of future demand by failing to communicate the inherent uncertainty in predictions. For grid operators, understanding the range of possible outcomes and their associated probabilities is essential for risk-informed decision-making, reserve procurement, and contingency planning. The transition from deterministic to probabilistic forecasting represents a significant methodological advancement in energy system ML [47].

Quantile regression directly estimates conditional quantiles of the target variable, providing prediction intervals at specified confidence levels without assuming a parametric error distribution. Quantile regression forests extend the random forest algorithm to produce quantile estimates by retaining the empirical distribution of training samples in each leaf node. Deep quantile regression networks employ pinball loss functions at multiple quantile levels to simultaneously estimate the 10th, 25th, 50th, 75th, and 90th percentiles of future demand, enabling the construction of flexible, non-parametric prediction intervals [48].

Bayesian Neural Networks (BNNs) place probability distributions over network weights rather than point estimates, enabling principled uncertainty quantification through posterior inference. The predictive distribution obtained by marginalizing over the weight posterior naturally decomposes into epistemic uncertainty (model uncertainty arising from limited training data) and aleatoric uncertainty (irreducible noise in the data). However, exact Bayesian inference in neural networks is intractable, necessitating approximate

methods such as variational inference (Bayes by Backprop) or Markov Chain Monte Carlo (MCMC) sampling, both of which significantly increase computational cost [49].

Monte Carlo (MC) Dropout, proposed by Gal and Ghahramani [50], provides a computationally efficient approximation to Bayesian inference by applying dropout at test time and interpreting the resulting stochastic forward passes as approximate samples from the posterior predictive distribution. For load forecasting, MC Dropout with 50-100 forward passes has been shown to produce well-calibrated prediction intervals at a fraction of the computational cost of full BNNs. Wang et al. [51] applied MC Dropout to an LSTM-based load forecasting model and demonstrated that the resulting 95% prediction intervals achieved coverage rates of 93-97% across different seasons and load levels, while point forecast accuracy remained comparable to deterministic LSTM baselines.

The importance of uncertainty quantification extends beyond academic interest to practical operational value. Grid operators in deregulated electricity markets can use probabilistic forecasts to optimize bidding strategies by accounting for the financial risks associated with forecast errors. System operators can calibrate operating reserve levels based on the predicted distribution of net load (demand minus renewable generation), potentially reducing reserve costs by 5-15% compared to deterministic approaches that apply fixed reserve margins. Distribution system operators can use prediction intervals to assess the probability of voltage limit violations and proactively deploy reactive power compensation [52].

4 | Demand-Side Management and Energy Efficiency

4.1 | Residential and Commercial Energy Optimization

DSM encompasses the portfolio of strategies, programs, and technologies designed to modify consumer electricity consumption patterns to improve system efficiency, reduce peak demand, and lower costs. ML techniques have fundamentally enhanced DSM capabilities by enabling intelligent automation, personalized optimization, and adaptive control of end-use loads at unprecedented granularity [53].

Smart Home Energy Management Systems (SHEMS) leverage ML to optimize the scheduling and control of flexible loads including HVAC systems, water heaters, dishwashers, washing machines, clothes dryers, and EV chargers while respecting consumer comfort constraints and exploiting time-varying electricity prices. DRL agents have demonstrated particular promise for SHEMS due to their ability to learn optimal control policies through trial-and-error interaction with the home environment without requiring explicit models of appliance dynamics or occupant behavior. Mocanu et al. [54] trained a Deep Q-Network agent to manage a simulated residential environment with five controllable loads under Real-Time Pricing (RTP), achieving electricity cost reductions of 18-25% compared to rule-based scheduling while maintaining thermal comfort within user-specified bounds. The DQN agent learned to pre-cool the home during off-peak price periods, defer flexible loads to overnight hours, and coordinate EV charging with anticipated solar generation.

NILM, also known as energy disaggregation, uses ML algorithms to decompose aggregate whole-house power measurements from a single smart meter into the contributions of individual appliances, eliminating the need for dedicated sub-meters on each device. Accurate disaggregation enables granular consumption feedback to consumers, identification of energy-wasteful behaviors, detection of malfunctioning appliances, and targeted efficiency recommendations. Deep learning has dramatically advanced NILM performance, with sequence-to-point (seq2point) models using CNNs and sequence-to-sequence models using Bi-LSTMs achieving state-of-the-art disaggregation accuracy. Kelly and Knottenbelt [55] demonstrated that a denoising autoencoder trained on UK-DALE data could disaggregate five target appliances with F1-scores exceeding 0.85 for high-power devices (kettle, microwave, dishwasher) and 0.70 for lower-power, continuously variable loads (refrigerator, washing machine). More recently, attention-based architectures and transformer models have further improved disaggregation accuracy by learning which temporal contexts are most informative for identifying appliance activation signatures [56].

User behavior modeling represents a critical yet often underexplored component of ML-driven DSM. Consumer energy consumption is fundamentally a socio-technical phenomenon shaped not only by physical building characteristics and weather conditions but also by occupant lifestyles, preferences, habits, and responsiveness to economic signals. Clustering techniques applied to smart meter data have revealed distinct consumer archetypes such as "morning peakers", "evening peakers", "flat consumers" and "night owls" that exhibit different load flexibility potentials and demand response propensities. Personalized energy recommendations generated by collaborative filtering algorithms, analogous to those used in e-commerce recommendation systems, have been shown to achieve 10-15% higher engagement rates compared to generic energy-saving tips by tailoring suggestions to individual consumption patterns and lifestyle characteristics [57].

4.2 | Industrial Energy Optimization

The industrial sector accounts for approximately 37% of global final energy consumption, making it a high-impact target for ML-driven optimization. Industrial energy management presents unique challenges including complex multi-stage production processes, stringent product quality constraints, variable raw material properties, and the interdependence of thermal, electrical, and mechanical energy flows. ML techniques are increasingly applied across several industrial optimization domains [58].

Predictive energy models in manufacturing leverage supervised learning algorithms to forecast the energy consumption of individual production machines, manufacturing cells, or entire production lines as a function of process parameters (feed rate, cutting speed, temperature setpoints), material properties, and production schedules. These predictive models enable energy-aware production planning, where manufacturing schedules are optimized to shift energy-intensive operations to periods of low electricity prices or high renewable generation availability while maintaining production throughput and quality targets. Gradient boosted tree models trained on historical production logs and energy meter data have demonstrated the ability to predict machine-level energy consumption with MAPE values of 3-7% for operations such as CNC machining, injection molding, and arc welding [59].

Real-time anomaly detection in industrial energy systems identifies deviations from expected energy consumption patterns that may indicate equipment degradation, process inefficiencies, compressed air leaks, steam trap failures, or fouled heat exchangers. Autoencoder-based anomaly detection, wherein a neural network is trained to reconstruct normal operating patterns and flags observations with high reconstruction error as anomalous, has been particularly effective for unsupervised fault detection in industrial settings where labeled fault data is scarce. Isolation forests and one-class SVMs provide alternative unsupervised approaches suitable for deployment on resource-constrained industrial edge devices [60].

The integration of on-site renewable generation with industrial production scheduling represents a multi-objective optimization challenge that ML is uniquely positioned to address. Industrial facilities with rooftop solar installations, on-site wind turbines, or co-generation plants must coordinate energy-intensive production activities with the availability of self-generated renewable electricity to minimize grid electricity purchases, reduce demand charges, and maximize the self-consumption ratio of renewable generation. Multi-objective RL agents have been demonstrated to simultaneously optimize production schedules, battery storage dispatch, and grid interaction to minimize total energy cost while meeting production deadlines and quality constraints [61].

4.3 | Dynamic Pricing and Incentive Mechanisms

Dynamic electricity pricing mechanisms, including Time-of-Use (TOU) tariffs, Critical Peak Pricing (CPP), and RTP, provide economic signals intended to incentivize consumers to shift flexible loads from peak to off-peak periods, thereby reducing the need for expensive peaking generation and transmission capacity. ML plays a dual role in dynamic pricing ecosystems: predicting future electricity prices to inform consumer scheduling decisions, and optimizing the design of pricing structures to maximize demand response effectiveness [62].

ML-based electricity price forecasting has become essential for consumers, aggregators, and retailers participating in deregulated electricity markets. Electricity spot prices exhibit complex dynamics including high volatility, heavy-tailed distributions, seasonality, mean reversion, and occasional extreme price spikes driven by supply shortages, transmission congestion, or strategic bidding behavior. Deep learning models, particularly LSTM networks and attention-based Transformers, have demonstrated superior performance in capturing these complex dynamics compared to traditional econometric models. Lago et al. [63] conducted a comprehensive benchmark of 27 forecasting methods across five European electricity markets, finding that deep neural networks achieved the lowest MAE in four of five markets, with ensemble methods combining deep learning and econometric models providing the best overall performance.

Game-theoretic approaches model the strategic interaction between utility providers and rational consumers in demand response programs. Stackelberg game formulations, wherein the utility acts as a leader setting prices and consumers act as followers optimizing their consumption in response, have been combined with ML-based demand models to derive pricing policies that maximize social welfare or utility profit while ensuring consumer participation constraints are satisfied. Multi-agent RL (MARL) frameworks extend this approach to settings with multiple heterogeneous consumers, enabling the discovery of equilibrium strategies that balance individual rationality with system-level objectives [64].

5 | Renewable Energy Integration and Grid Stability

5.1 | Solar and Wind Power Forecasting

The accurate forecasting of solar and wind power generation is a prerequisite for the reliable and economic integration of VRES into the power grid. Forecast errors directly translate to imbalance costs, increased reserve requirements, curtailment of renewable generation, and potential reliability risks. The challenge of VRES forecasting stems from the stochastic nature of meteorological processes, the spatial variability of weather conditions across distributed renewable installations, and the nonlinear mapping between weather variables and power output [65].

Solar power forecasting methods can be categorized by forecast horizon. For very short-term horizons (seconds to minutes), sky camera-based cloud tracking combined with CNN-based image classification provides nowcasting capabilities that support real-time grid balancing and inverter control. For intra-day horizons (one to six hours), satellite-derived cloud motion vectors and statistical time-series models are commonly employed. For day-ahead and longer horizons, NWP models provide the primary information source, with ML post-processing techniques correcting systematic biases and improving spatial and temporal resolution [66].

PINNs have emerged as a promising approach for renewable forecasting by embedding known physical relationships into the ML model structure. For solar forecasting, PINNs can incorporate the clear-sky irradiance model (which describes the theoretical maximum irradiance as a function of solar position, atmospheric constituents, and elevation) as a baseline, with the neural network learning the deviation caused by cloud cover, aerosols, and other atmospheric perturbations. This physics-informed approach improves generalization to unseen weather conditions and reduces the training data requirements by leveraging domain knowledge to constrain the hypothesis space [67].

Hybrid models combining NWP outputs with deep learning post-processing have achieved state-of-the-art performance for day-ahead wind power forecasting. A typical pipeline ingests ensemble NWP forecasts of wind speed, wind direction, temperature, and pressure at multiple atmospheric levels, applies spatial interpolation to the turbine locations, and trains a deep neural network to map the NWP features to turbine-level power output while correcting for terrain effects, wake interactions, and NWP biases. Wang et al. [68] demonstrated that a CNN-LSTM model post-processing ECMWF ensemble forecasts achieved a normalized RMSE of 8.3% for day-ahead wind farm power prediction, representing a 31% improvement over the raw NWP persistence ensemble baseline.

Spatial-temporal correlation modeling across geographically distributed renewable assets provides additional forecasting value by exploiting the propagation patterns of weather systems. Graph neural networks (GNNs), which operate on graph-structured data where nodes represent individual solar or wind installations and edges encode spatial relationships, have demonstrated particular promise for this application. The message-passing mechanism of GNNs enables each node to aggregate information from neighboring nodes, capturing the spatial correlation of generation patterns that arise from shared weather conditions and sequential cloud shadow or wind front propagation [69].

5.2 | Energy Storage Management

ESS, particularly lithium-ion batteries, are critical enabling technologies for smart grids, providing services including peak shaving, frequency regulation, renewable generation smoothing, and backup power. ML plays increasingly important roles in both the operational optimization and health management of energy storage assets [70].

Battery degradation modeling using ML addresses the need to predict the Remaining Useful Life (RUL) and capacity fade trajectory of battery systems, which is essential for warranty management, replacement planning, and economic viability assessment. Battery degradation is a complex electrochemical process influenced by cycling depth, C-rate, temperature, State of Charge (SOC) range, and calendar aging, making physics-based degradation models computationally expensive and often insufficiently accurate for real-world operating conditions. ML-based approaches, trained on accelerated aging test data, can learn empirical degradation models that capture the nonlinear interactions among stress factors. Severson et al. [71] demonstrated that an ML model using features extracted from the first 100 charge-discharge cycles could predict battery cycle life with 9.1% mean percentage error, enabling early identification of cells destined for premature failure.

Optimal charge/discharge scheduling of BESS is naturally formulated as a sequential decision-making problem amenable to DRL. The DRL agent observes the current state (SOC, electricity price, net load forecast, time of day, battery health indicators), selects charge/discharge power setpoints, and receives rewards based on energy cost savings, revenue from ancillary service provision, and penalties for battery degradation. Cao et al. [72] trained a PPO agent for battery arbitrage in a real-time electricity market and demonstrated annualized revenue increases of 12-18% compared to rule-based and mixed-integer programming approaches, with the additional advantage of implicit degradation-aware operation without requiring an explicit battery aging model.

Vehicle-to-Grid (V2G) coordination and EV fleet management represent emerging applications where ML techniques are essential for managing the complexity arising from heterogeneous vehicle batteries, stochastic arrival and departure patterns, diverse consumer preferences, and the dual objective of satisfying mobility needs while providing grid services. MARL frameworks, where each EV is represented by an autonomous agent that learns cooperative charging/discharging strategies, have shown promise for scalable V2G coordination. Hydrogen storage optimization, though less mature, leverages similar ML-based scheduling approaches for electrolyzer and fuel cell dispatch in power-to-gas-to-power configurations [73].

5.3 | Grid Stability and Frequency Regulation

The displacement of synchronous generators by inverter-based renewable resources fundamentally alters the dynamic characteristics of power systems, reducing rotational inertia and weakening the inherent frequency response that has historically maintained system stability. ML techniques are being developed to address the emerging challenges of grid stability assessment and frequency regulation in low-inertia power systems [74].

ML-based voltage stability assessment replaces computationally intensive power flow calculations and continuation power flow analyses with trained surrogate models that can rapidly evaluate stability margins from real-time measurements. Deep neural networks trained on thousands of simulated operating conditions can classify system states as stable or unstable, estimate voltage stability margins, and identify critical

contingencies in milliseconds rather than the minutes required by conventional numerical methods. This speed advantage is essential for real-time situational awareness in control centers [75].

Inertia estimation and frequency response prediction using ML address the challenge of quantifying system inertia in real-time, which is essential for determining frequency containment reserve requirements and activating emergency load shedding schemes. Data-driven approaches utilizing PMU measurements of frequency, Rate of Change of Frequency (RoCoF), and power imbalance during disturbance events can train regression models that estimate the effective system inertia with accuracy sufficient for operational decision-making [76].

Wide-area monitoring and control using PMU data and ML analytics enables the detection of inter-area oscillations, voltage instability, and impending cascading failures in real-time. CNNs applied to PMU data matrices (where rows represent PMU channels and columns represent time steps) have demonstrated the ability to detect and classify inter-area oscillation modes with accuracy exceeding 95%, enabling early warning of potential instability. Islanding detection, the identification of distribution network sections that have become electrically isolated from the main grid, is critical for safety and is addressed using classification algorithms (SVM, random forest, neural networks) trained on features extracted from voltage, frequency, and harmonic measurements at the point of common coupling [77].

6 | Reinforcement Learning for Real-Time Grid Control

6.1 | Reinforcement Learning Formulation for Smart Grid Problems

The application of RL to smart grid control requires careful formulation of the problem as a Markov Decision Process (MDP) or its generalizations. The state space, action space, reward function, and transition dynamics must be designed to faithfully represent the physical system while remaining computationally tractable for RL algorithms [78].

State space design for smart grid RL problems typically includes the current time step, historical load and generation measurements, weather forecasts, electricity price signals, equipment status indicators, and relevant physical quantities such as bus voltages, line flows, generator outputs, and battery SOC levels. The high dimensionality of realistic grid state representations potentially encompassing thousands of buses, generators, and loads necessitates state compression techniques including principal component analysis, autoencoders, and state aggregation heuristics. Partial observability, arising from limited sensor deployment, communication delays, and measurement noise, motivates the use of RNNs (LSTM or GRU) within the RL policy to maintain internal state representations that integrate information over time [79].

Action space design presents distinct challenges depending on whether the control problem involves discrete decisions (e.g., switching, breaker operations, tap changer positions) or continuous setpoints (e.g., generator power output, voltage reference, battery charge/discharge rate). Hybrid action spaces combining discrete and continuous components are common in realistic grid control problems. Parameterized action space methods and branching DQN architectures have been proposed to handle such hybrid settings efficiently [80].

Reward function design is arguably the most critical and domain-specific aspect of RL formulation for smart grids. The reward must balance multiple, potentially conflicting objectives including economic cost minimization, voltage regulation quality, frequency stability, equipment life preservation, environmental impact reduction, and consumer satisfaction. Poorly designed reward functions can lead to unintended and potentially dangerous behaviors, such as an agent that minimizes cost by consistently operating equipment at thermal limits. Multi-objective reward formulations using weighted sums, lexicographic ordering, or Pareto-based approaches have been explored, though the selection of appropriate objective weights often requires domain expertise and iterative refinement [7].

Model-based RL approaches learn an explicit model of the environment's transition dynamics and use this model for planning, potentially achieving much greater sample efficiency than model-free methods. For smart

grid applications, the environment model can be initialized with known physical models (power flow equations, generator dynamics, battery electrochemistry) and refined with observational data, creating a hybrid physics-ML simulator that enables rapid, low-cost exploration of control strategies before deployment on the physical grid. World models and Dreamer-style architectures, which learn latent state representations and transition dynamics jointly, have shown promising results for complex energy system control tasks [81].

6.2 | Multi-Agent Reinforcement Learning

The inherently distributed nature of modern smart grids comprising thousands of autonomous prosumers, distributed generators, storage devices, and controllable loads motivates the application of MARL for decentralized coordination. In MARL frameworks, each DER is represented by an autonomous agent that observes local information, takes local actions, and learns a policy that contributes to system-level objectives through cooperation or competition with other agents [82].

Cooperative MARL frameworks address problems where all agents share a common objective, such as minimizing total system cost, maximizing renewable self-consumption, or maintaining network voltage within acceptable bounds. Centralized Training with Decentralized Execution (CTDE) architectures, such as QMIX and MAPPO, enable agents to learn coordinated policies during training while executing independently based on local observations during deployment, avoiding the prohibitive communication requirements of fully centralized control. QMIX factorizes the joint action-value function into individual agent value functions subject to a monotonicity constraint, enabling tractable optimization while preserving coordination capabilities [83].

Peer-to-Peer (P2P) energy trading represents a competitive MARL application where prosumers consumers with distributed generation and storage negotiate bilateral energy trades to exploit local supply-demand complementarities, reducing reliance on grid imports and associated network charges. Each prosumer agent learns a bidding strategy that maximizes individual profit while accounting for the strategic behavior of trading partners. Mean-field game approximations, which replace the interactions with individual agents by interactions with the aggregate population behavior, have been employed to scale P2P trading algorithms to thousands of participants [84].

Scalability remains the primary challenge for MARL in smart grid applications. As the number of agents increases, the joint action space grows exponentially, the learning landscape becomes non-stationary (since each agent's environment is influenced by the evolving policies of other agents), and communication overhead between agents and a central coordinator can become prohibitive. Hierarchical MARL architectures, where agents are organized into groups managed by higher-level coordinators, and attention-based communication mechanisms that enable selective information sharing, have been proposed to address these scalability challenges [85].

6.3 | Safe and Constrained Reinforcement Learning

The deployment of RL agents for controlling critical energy infrastructure demands rigorous safety guarantees that standard RL algorithms, designed to maximize cumulative reward without explicit constraint enforcement, cannot provide. An RL agent controlling generator dispatch or voltage regulators must never violate physical constraints (thermal line limits, voltage magnitude bounds, generator ramp rate limits) or operational constraints (minimum up/down times, reserve requirements) even during the exploration phase of learning, as violations can damage equipment, endanger personnel, or trigger cascading blackouts [86].

Constrained Markov Decision Processes (CMDPs) extend the standard MDP framework by introducing auxiliary cost functions and associated budget constraints that the agent must satisfy in expectation over the learning trajectory. Lagrangian relaxation methods convert the constrained optimization problem into an unconstrained one by introducing dual variables (Lagrange multipliers) that penalize constraint violations, and alternately updating the policy and the dual variables using primal-dual optimization. The Constrained Policy

Optimization (CPO) algorithm provides approximate constraint satisfaction guarantees at each policy update step by solving a trust-region optimization problem with linearized constraints [87].

Sim-to-real transfer addresses the fundamental challenge that RL agents trained in simulation may perform poorly when deployed on physical power systems due to discrepancies between the simulated and real environments (the "reality gap"). Domain randomization, where simulation parameters (line impedances, load models, generator characteristics) are randomly varied during training to produce a robust policy that generalizes across a distribution of possible environments, has shown effectiveness for grid control tasks. System identification techniques that progressively calibrate the simulation environment to match real-world observations, combined with fine-tuning of pre-trained policies on limited real-world data, provide a practical pathway for safe RL deployment in grid applications [88].

Safety layers and action projection methods provide an additional safeguard by projecting the RL agent's proposed actions onto the feasible set defined by physical and operational constraints before execution. A trained neural network-based safety filter can verify in real-time whether a proposed action is feasible and, if not, find the nearest feasible action using quadratic programming or barrier function methods. This approach decouples the learning of optimal behavior (handled by the RL algorithm) from the enforcement of safety constraints (handled by the projection layer), simplifying both components [89].

7 | Privacy, Security, and Federated Learning

7.1 | Privacy Concerns in Smart Grid Machine Learning

The deployment of ML techniques in smart grids, while offering substantial operational benefits, raises significant privacy concerns arising from the granularity and sensitivity of the data required to train effective models. High-resolution smart meter data, typically collected at 15-minute or finer intervals, can reveal detailed information about household activities, occupancy patterns, sleep schedules, appliance ownership, and lifestyle habits. Research has demonstrated that NILM algorithms can infer when specific appliances are operated, potentially revealing sensitive activities such as medical device usage, home occupancy during vacation periods, or irregular sleep patterns indicative of health conditions [90].

Regulatory frameworks, most notably the European Union's General Data Protection Regulation (GDPR), impose strict requirements on the collection, processing, storage, and sharing of personal data, including energy consumption data that can be linked to identifiable individuals. Under GDPR, smart meter data constitutes personal data, and its use for ML model training requires a lawful basis such as explicit consent, legitimate interest, or contractual necessity. The right to data minimization requires that only the minimum amount of data necessary for the specified purpose be collected and processed, potentially conflicting with the data-hungry nature of deep learning models. The right to explanation requires that automated decisions affecting individuals be explainable, posing challenges for opaque deep learning models used in dynamic pricing or credit assessment based on energy consumption patterns [91].

Differential Privacy (DP) provides a mathematically rigorous framework for quantifying and bounding the privacy loss associated with data analysis. A randomized algorithm satisfies ϵ -DP if the output distribution changes by at most a multiplicative factor of e^ϵ when any single individual's data is added or removed from the input dataset. In the smart grid context, differentially private mechanisms add calibrated noise to aggregated consumption statistics, model gradients, or query responses to prevent the inference of individual consumption patterns. The fundamental tension between privacy and utility stronger privacy guarantees require larger noise additions that degrade model accuracy is characterized by the privacy budget ϵ , with typical practical deployments using ϵ values between 1 and 10. Eibl and Engel [92] demonstrated that differentially private load profiles with $\epsilon = 1$ could support aggregate demand forecasting with less than 3% MAPE degradation while providing strong individual privacy guarantees.

7.2 | Federated Learning for Distributed Smart Grids

FL has emerged as the most promising paradigm for privacy-preserving collaborative ML in smart grid systems, enabling multiple parties utilities, consumers, or grid operators to jointly train shared ML models without exchanging raw data. In the canonical FL framework (FedAvg), each participant trains a local model on its private data, shares only the model parameter updates (gradients) with a central aggregation server, which computes a weighted average of the updates and distributes the updated global model back to participants. This process iterates until convergence, with raw data never leaving the originating device or organization [93].

Horizontal FL, where participants share the same feature space but hold different data samples, is the most natural configuration for smart grid applications. A typical deployment involves hundreds or thousands of smart meters, each holding years of consumption data for a single household, collaboratively training a shared load forecasting model without revealing individual consumption patterns to the utility or other consumers. Taik and Cherkaoui [94] demonstrated that a federated LSTM model for residential load forecasting achieved accuracy within 2-5% of a centrally trained model (with all data pooled) while providing strong privacy guarantees, representing a favorable privacy-utility tradeoff.

Vertical FL, where participants hold different features for the same set of entities, is applicable when multiple organizations contribute complementary data sources for joint modeling. For example, a utility holding consumption data and a weather service holding meteorological data could collaboratively train a forecasting model without either party revealing its proprietary data to the other. Secure computation protocols, including homomorphic encryption and secure multi-party computation, protect the confidentiality of intermediate computations during the vertical FL training process [95].

Several technical challenges complicate the deployment of FL in smart grid environments. Statistical heterogeneity (non-IID data), arising from the diverse consumption patterns of different households, geographical regions, and building types, degrades the performance of standard FedAvg aggregation. Personalization techniques, including local fine-tuning, meta-learning (per-FedAvg), and mixture-of-experts approaches, have been proposed to adapt the global model to the local data distribution of each participant while retaining the benefits of collaborative training. Communication efficiency is critical given the bandwidth-constrained communication channels in smart grid networks; gradient compression techniques (sparsification, quantization, and sketching) reduce the communication payload by orders of magnitude with minimal accuracy loss. Byzantine fault tolerance mechanisms protect against malicious participants who submit poisoned model updates aimed at degrading global model performance or embedding backdoors [96].

7.3 | Cybersecurity and Adversarial Robustness

The increasing reliance of smart grid operations on ML-based decision-making introduces novel cybersecurity vulnerabilities that extend beyond traditional IT security concerns. Adversarial ML, which studies the susceptibility of ML models to carefully crafted malicious inputs, has revealed that deep learning models used for load forecasting, state estimation, and fault detection can be manipulated by adversaries with knowledge of the model architecture or training data [97].

False Data Injection Attacks (FDIAs) on ML-based state estimation represent a critical threat, where an adversary compromises a subset of sensor measurements and injects carefully constructed false readings that bypass conventional bad data detection mechanisms while causing the state estimator to produce erroneous results. ML-based state estimators, while more accurate than traditional weighted least squares methods under normal conditions, can be more susceptible to adversarial perturbations due to their sensitivity to input distribution shifts. Defense strategies include adversarial training (augmenting the training dataset with adversarial examples), robust optimization formulations that minimize worst-case loss over perturbation sets, and ensemble methods that combine multiple diverse estimators to increase the difficulty of simultaneous evasion [98].

Adversarial examples targeting load forecasting models can cause systematic prediction biases that lead to suboptimal dispatch decisions, increased operating costs, or reliability violations. Chen et al. [99] demonstrated that imperceptible perturbations to weather input features could shift day-ahead load forecasts by up to 15% without triggering anomaly detection alarms, potentially causing millions of dollars in economic losses through inefficient generation scheduling. Certified robustness methods, including randomized smoothing and interval bound propagation, provide provable guarantees on the maximum prediction change under bounded input perturbations, though computational costs and conservatism remain practical limitations.

Anomaly-based Intrusion Detection Systems (IDS) using deep learning offer a complementary defense by monitoring network traffic, sensor data streams, and system logs for patterns indicative of cyberattacks. Autoencoders and VAEs trained on normal grid operating data can detect anomalous patterns with high sensitivity, while recurrent models (LSTM, GRU) capture the temporal structure of attack sequences. The challenge of extremely low false positive rates required for operational deployment grid operators cannot investigate more than a handful of alerts per day necessitates sophisticated threshold selection and alert prioritization mechanisms [100].

8| Emerging Trends and Future Directions

8.1| Edge Artificial Intelligence and On-Device Intelligence

The deployment of ML models on edge devices smart meters, IoT sensors, protection relays, and distribution automation controllers promises to transform smart grid operations by enabling real-time, low-latency inference without dependence on cloud connectivity. Edge AI is particularly compelling for applications requiring sub-second response times (protection, frequency regulation), operating in bandwidth-constrained environments (remote substations, rural feeders), or processing privacy-sensitive data that should not leave the device [101].

Model compression techniques reduce the computational and memory requirements of deep learning models to levels compatible with resource-constrained edge hardware. Structured pruning removes entire filters, channels, or attention heads from neural networks based on importance criteria (magnitude, gradient, or sensitivity), achieving 50-90% parameter reduction with less than 2% accuracy degradation for many forecasting tasks. Quantization reduces the numerical precision of weights and activations from 32-bit floating point to 8-bit integers or lower, reducing model size by 4× and accelerating inference on hardware with integer arithmetic support. Knowledge distillation trains a compact "student" model to mimic the output distribution of a larger "teacher" model, transferring the teacher's learned knowledge into a deployable model architecture. Combined application of these techniques has demonstrated the feasibility of deploying LSTM-based load forecasting models on microcontrollers with less than 256 KB of RAM [102].

8.2| Large Language Models and Foundation Models

The emergence of LLMs and foundation models represents a paradigm shift in artificial intelligence that is beginning to influence energy systems research. While LLMs are primarily designed for natural language processing, their potential applications in smart grid optimization extend across several dimensions that leverage their general reasoning capabilities and vast encoded knowledge [103].

LLM-assisted energy system design and optimization utilizes the natural language understanding and code generation capabilities of models such as GPT-4, Claude, and open-source alternatives to accelerate the development of energy system models, optimization formulations, and control algorithms. Energy engineers can describe optimization problems in natural language and receive mathematical formulations, Python code implementations, and solution recommendations, significantly reducing development time and democratizing access to advanced optimization techniques. Fine-tuned LLMs can also serve as intelligent assistants for grid

operators, translating alarm sequences and system status information into natural language situation reports and recommended response actions [104].

Foundation models for time-series data, pre-trained on massive corpora of time-series data from diverse domains, are emerging as a new paradigm for energy forecasting. TimesFM, Chronos, and Lag-Llama represent early examples of time-series foundation models that demonstrate zero-shot and few-shot forecasting capabilities competitive with task-specific models trained from scratch. For smart grid applications, these foundation models offer the potential to provide accurate forecasts for new deployments, uncommon load profiles, and data-scarce scenarios without requiring extensive labeled training data. The adaptation of these models to the specific characteristics of energy data, including strong periodicity, calendar effects, and physical constraints through domain-specific fine-tuning and prompt engineering remains an active research area [105].

8.3 | Digital Twins and Simulation-Based Optimization

Digital Twins (DT), high-fidelity virtual replicas of physical power systems that are continuously synchronized with real-time operational data, provide a powerful platform for ML model development, validation, and deployment. A digital twin of a distribution network, for example, integrates the network topology, equipment parameters, load models, DER models, and weather data to create a virtual environment that faithfully reproduces the electrical behavior of the physical system under various operating conditions and contingencies [106].

Physics-informed DT combines first-principles physical models (power flow equations, electromagnetic transient models, thermal models) with data-driven components (ML-based load models, renewable generation forecasters, equipment degradation estimators) to achieve higher fidelity than either approach alone. The digital twin enables ML researchers to generate unlimited synthetic training data under controlled conditions, including rare events (faults, extreme weather, cascading failures) that are insufficiently represented in historical operational data. This synthetic data generation capability is particularly valuable for training RL agents, where exploration of dangerous states in the physical system is unacceptable [107].

Generative Adversarial Networks (GANs) offer an alternative approach to synthetic data generation for smart grid ML, learning to generate realistic consumption profiles, renewable generation traces, and fault signatures from training data. Conditional GANs can generate synthetic data conditioned on specified attributes (season, building type, occupancy level), enabling targeted augmentation of underrepresented scenarios in the training dataset. The quality of GAN-generated synthetic energy data has been shown to be sufficient for training downstream forecasting and classification models with less than 5% accuracy degradation compared to models trained on real data, while providing DP guarantees by avoiding the direct use of sensitive real consumption data [108].

8.4 | Explainable Artificial Intelligence for Energy Systems

The adoption of ML models for safety-critical smart grid applications including protection system settings, emergency load shedding, and investment planning requires that model predictions be interpretable, auditable, and trustworthy. The "black box" nature of deep learning models, which make predictions through millions of parameters without providing human-understandable explanations, represents a significant barrier to acceptance by grid operators, regulators, and consumers [109].

Post-hoc explainability methods generate explanations for the predictions of pre-trained models without modifying the model architecture. Shapley Additive Explanations (SHAP) values, grounded in cooperative game theory, attribute the prediction of each instance to the contribution of each input feature, satisfying desirable properties including local accuracy, missingness, and consistency. Applied to load forecasting models, SHAP analysis can reveal that temperature is the dominant driver during summer peaks, while day-of-week and time-of-day features dominate during mild weather periods, providing actionable insights for grid operators. Local Interpretable Model-agnostic Explanations (LIME) constructs locally faithful linear

approximations around individual predictions, generating sparse, human-readable explanations that highlight the most influential features for each specific forecast [110].

Attention-based interpretability, inherent in Transformer architectures, provides built-in explanations through the attention weight matrices, which reveal which input time steps and features the model focuses on when generating predictions. The TFT explicitly separates static and time-varying feature importance, producing interpretable attention patterns that can be directly communicated to grid operators. However, recent research has questioned whether attention weights truly reflect model reasoning or merely correlate with feature importance, motivating the development of more rigorous attention-based explanation methods [111].

Regulatory requirements for explainable automated decisions are becoming increasingly stringent. The EU AI Act classifies AI systems used in critical infrastructure management, including energy systems, as "high-risk" requiring risk assessments, human oversight, transparency documentation, and the ability to explain individual decisions to affected parties. Compliance with these regulatory requirements necessitates not only technical explainability but also organizational processes for model governance, documentation, monitoring, and periodic re-evaluation [112].

8.5 | Open Challenges and Research Gaps

Despite the substantial progress reviewed in preceding sections, several critical challenges must be addressed to enable the widespread deployment of ML-driven optimization in production smart grid environments.

Data standardization and interoperability remain fragmented across the global energy industry. Different utilities, equipment manufacturers, and grid operators employ incompatible data formats, naming conventions, measurement resolutions, and communication protocols, hindering the development of transferable ML models and benchmark comparisons across studies. Industry-wide adoption of standards such as the Common Information Model (CIM), IEC 61850, and OpenADR would facilitate data sharing, model portability, and reproducible research [113].

Scalability of ML solutions from controlled pilot projects to full grid-scale deployment presents engineering challenges that are frequently underappreciated in academic research. Models demonstrating excellent performance on benchmark datasets with thousands of data points may encounter difficulties with the terabytes of streaming data generated by millions of smart meters. Real-time inference latency constraints, model retraining frequencies, data pipeline reliability, and graceful degradation under data quality issues require production engineering practices that extend beyond algorithm development [114].

Multi-objective optimization that simultaneously balances economic cost, carbon emissions, reliability, equity, and resilience represents a fundamentally harder problem than single-objective optimization and remains insufficiently addressed by current ML approaches. The tradeoffs between these objectives are often non-convex, non-stationary, and politically contested, requiring transparent articulation of preference structures and participatory decision-making processes that involve diverse stakeholders [115].

Cross-domain transfer learning between different grid topologies, voltage levels, climate regions, and regulatory environments offers the potential to dramatically accelerate ML deployment but is limited by the current understanding of domain similarity measures and transfer conditions for energy systems. Negative transfer, where knowledge from the source domain degrades performance on the target domain, remains poorly characterized and difficult to predict a priori [116].

Table 2. Summary of key research gaps and proposed future directions.

Challenge Area	Current Limitation	Promising Research Direction	Expected Impact
Data scarcity	Insufficient labeled data for rare events	Synthetic data via GANs; transfer learning; foundation models	Faster deployment in new environments
Model interpretability	Black-box models hinder trust and regulatory compliance	SHAP, attention mechanisms, inherently interpretable models	Regulatory approval, operator adoption
Scalability	Pilot-to-production gap	Edge AI, model compression, streaming ML pipelines	Real-time grid-scale deployment
Privacy	Smart meter data reveals sensitive information	FL, DP, secure aggregation	GDPR compliance, consumer trust
Safety	RL agents may violate physical constraints	Constrained RL, safety layers, sim-to-real transfer	Safe autonomous grid control
Adversarial robustness	ML models vulnerable to adversarial attacks	Adversarial training, certified robustness, ensemble defenses	Resilient critical infrastructure
Multi-objective optimization	Single-objective formulations dominate	Pareto-based RL, multi-objective Bayesian optimization	Balanced cost-emission-reliability tradeoffs
Standardization	Incompatible data formats across utilities	CIM adoption, open benchmarks, federated data spaces	Transferable models, reproducible research

9 | Conclusion

This comprehensive review has systematically examined the rapidly evolving landscape of ML-driven optimization for energy consumption in smart grids, synthesizing over 100 studies published between 2018 and 2025 across the domains of load forecasting, DSM, renewable energy integration, grid stability, real-time control, privacy preservation, and emerging methodological frontiers.

The evidence reviewed demonstrates that ML techniques have achieved transformative performance improvements across virtually every aspect of smart grid optimization. In load forecasting, deep learning architectures, particularly LSTM networks, CNN-LSTM hybrids, and TFT, have reduced prediction errors by 20-40% compared to traditional statistical methods, with attention-based models providing the additional benefit of interpretable feature importance attribution. For DSM, DRL agents have demonstrated 15-25% electricity cost reductions in residential settings while maintaining consumer comfort, and NILM algorithms powered by deep learning enable granular appliance-level energy feedback without intrusive sub-metering. Renewable energy forecasting has been substantially improved through hybrid approaches combining NWP with deep learning post-processing, achieving day-ahead wind and solar power prediction accuracies that enable higher renewable penetration levels without compromising grid reliability.

RL has emerged as the most promising paradigm for real-time adaptive grid control, with multi-agent frameworks enabling scalable coordination of DERs and P2P energy trading. However, the deployment of RL agents in safety-critical grid infrastructure requires constrained RL formulations, safety layers, and rigorous sim-to-real transfer methodologies that remain active research areas. FL has established itself as the leading framework for privacy-preserving collaborative model training, enabling utilities and consumers to benefit from shared intelligence without compromising data confidentiality, though challenges related to non-IID data distributions and communication efficiency persist.

Several critical barriers must be overcome before ML-driven optimization can achieve widespread adoption in production smart grid environments. Data quality, standardization, and interoperability remain fragmented, hindering model transferability and reproducible benchmarking. Model interpretability and explainability are essential for operator trust, regulatory compliance under emerging legislation such as the EU AI Act, and safe human-in-the-loop decision-making. Scalability from controlled academic experiments to grid-scale deployment with millions of endpoints requires production engineering capabilities that extend beyond algorithmic innovation. Cybersecurity and adversarial robustness must be integrated into ML model design from the outset, given the critical nature of energy infrastructure.

Looking forward, several emerging trends hold transformative potential. Edge AI and TinyML will enable real-time, on-device intelligence across the smart grid's sensor network, reducing latency and communication burden. Foundation models for time-series data promise zero-shot and few-shot forecasting capabilities that could dramatically accelerate deployment in data-scarce environments. DT will provide high-fidelity simulation environments for safe RL training and ML model validation. Explainable AI methods will bridge the trust gap between ML capabilities and operational acceptance.

The realization of ML-driven smart grid optimization at scale will require sustained interdisciplinary collaboration among ML researchers, power systems engineers, cybersecurity experts, regulatory authorities, and policymakers. The technical capabilities reviewed in this paper are maturing rapidly, but their translation into operational impact depends equally on organizational readiness, regulatory frameworks, workforce development, and public acceptance. The convergence of these technical and socio-institutional factors will determine the pace at which ML transforms the global energy system toward a more efficient, reliable, sustainable, and equitable future.

Authors' Contributions

All aspects of the research and manuscript preparation were carried out by the author. The author has read and approved the final version of the manuscript.

Funding

Not applicable.

Data Availability

All data supporting the reported findings in this research paper are provided within the manuscript.

Conflict of Interest

The author declares that they do not have any conflict of interest.

Consent for Publication

The author confirms consent for the publication of this work

Ethics Approval and Consent to Participate

This article does not contain any studies with human participants performed by the author.

References

- [1] International Energy Agency (IEA). (2022). *World energy outlook 2022*. Paris, france: International energy agency (IEA), 17. <https://www.iea.org/reports/world-energy-outlook-2022>
- [2] Lasseter, R. H., & Paigi, P. (2004). Microgrid: A conceptual solution. *2004 IEEE 35th annual power electronics specialists conference* (Vol. 6, pp. 4285–4290). IEEE. <https://doi.org/10.1109/PESC.2004.1354758>
- [3] Fang, X., Misra, S., Xue, G., & Yang, D. (2011). Smart grid – The new and improved power grid: A survey. *IEEE communications surveys & tutorials*, 14(4), 944–980. <https://doi.org/10.1109/SURV.2011.101911.00087>
- [4] Gungor, V. C., Sahin, D., Kocak, T., Ergut, S., Buccella, C., Cecati, C., & Hancke, G. P. (2011). Smart grid technologies: Communication technologies and standards. *IEEE transactions on industrial informatics*, 7(4), 529–539. <https://doi.org/10.1109/TII.2011.2166794>
- [5] Rawat, D. B., & Bajracharya, C. (2015). Cyber security for smart grid systems: Status, challenges and perspectives. *SoutheastCon 2015* (pp. 1–6). Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/SECON.2015.7132891>

- [6] Yan, Y., Qian, Y., Sharif, H., & Tipper, D. (2012). A survey on cyber security for smart grid communications. *IEEE communications surveys & tutorials*, 14(4), 998–1010. <https://doi.org/10.1109/SURV.2012.010912.00035>
- [7] Frank, S., Steponavice, I., & Rebennack, S. (2012). Optimal power flow: A bibliographic survey I: Formulations and deterministic methods. *Energy systems*, 3(3), 221–258. <https://doi.org/10.1109/JPROC.2020.2988715>
- [8] Deng, R., Yang, Z., Chow, M. Y., & Chen, J. (2015). A survey on demand response in smart grids: Mathematical models and approaches. *IEEE transactions on industrial informatics*, 11(3), 570–582. <https://doi.org/10.1109/TII.2015.2414719>
- [9] Siano, P. (2014). Demand response and smart grids — A survey. *Renewable and sustainable energy reviews*, 30, 461–478. <https://doi.org/10.1016/j.rser.2013.10.022>
- [10] Phadke, A. G., & Thorp, J. S. (2008). *Synchronized phasor measurements and their applications*. Springer. <https://doi.org/10.1007/978-3-319-50584-8>
- [11] Ackermann, T., Andersson, G., & Söder, L. (2001). Distributed generation: A definition. *Electric power systems research*, 57(3), 195–204. [https://doi.org/10.1016/S0378-7796\(01\)00101-8](https://doi.org/10.1016/S0378-7796(01)00101-8)
- [12] Alahakoon, D., & Yu, X. (2015). Smart electricity meter data intelligence for future energy systems: A survey. *IEEE transactions on industrial informatics*, 12(1), 425–436. <https://doi.org/10.1109/TII.2015.2414355>
- [13] Wang, Y., Chen, Q., Hong, T., & Kang, C. (2018). Review of smart meter data analytics: Applications, methodologies, and challenges. *IEEE transactions on smart grid*, 10(3), 3125–3148. <https://doi.org/10.1109/TSG.2018.2818167>
- [14] Wang, W., Xu, Y., & Khanna, M. (2011). A survey on the communication architectures in smart grid. *Computer networks*, 55(15), 3604–3629. <https://doi.org/10.1016/j.comnet.2011.07.010>
- [15] Smola, A. J., & Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and computing*, 14(3), 199–222. <https://doi.org/10.1023/B:STCO.0000035301.49549.88>
- [16] Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [17] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785–794). Association for Computing Machinery (ACM). <https://doi.org/10.1145/2939672.2939785>
- [18] Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3), 273–297. <https://doi.org/10.1007/BF00994018>
- [19] Fahim, M., & Sillitti, A. (2019). Anomaly detection, analysis and prediction techniques in iot environment: A systematic literature review. *IEEE access*, 7, 81664–81681. <https://doi.org/10.1109/ACCESS.2019.2921912>
- [20] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [21] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [22] Ashish, V. (2017). Attention is all you need. *Advances in neural information processing systems*, 30, I. <https://cit.nii.ac.jp/crid/1370849946232757637>
- [23] Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International journal of forecasting*, 37(4), 1748–1764. <https://doi.org/10.1016/j.ijforecast.2021.03.012>
- [24] Kingma, D. P., & Welling, M. (2013). *Auto-encoding variational bayes*. <https://doi.org/10.48550/arXiv.1312.6114>
- [25] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- [26] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). *Proximal policy optimization algorithms*. <https://doi.org/10.48550/arXiv.1707.06347>
- [27] Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International conference on machine learning* (pp. 1861–1870). Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v80/haarnoja18b/haarnoja18b.pdf>

- [28] Kim, S., & Kang, M. (2019). *Financial series prediction using attention LSTM*. <https://doi.org/10.48550/arXiv.1902.10877>
- [29] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of computational physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [30] Hong, T., Pinson, P., & Fan, S. (2014). Global energy forecasting competition 2012. *International journal of forecasting*, 30(2), 357–363. <https://doi.org/10.1016/j.ijforecast.2013.07.001>
- [31] Davis, J., & Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. *Proceedings of the 23rd international conference on machine learning* (pp. 233–240). Association for Computing Machinery (ACM). <https://doi.org/10.1145/1143844.1143874>
- [32] Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE transactions on industrial informatics*, 7(3), 381–388. <https://doi.org/10.1109/TII.2011.2158841>
- [33] Kolter, J. Z., & Johnson, M. J. (2011). REDD: A public data set for energy disaggregation research. *Workshop on data mining applications in sustainability (SIGKDD)*, San Diego, CA (pp. 59–62). Association for Computing Machinery (ACM). https://www.researchgate.net/publication/266597071_REDD_A_Public_Data_Set_for_Energy_Disaggregation_Research
- [34] Kelly, J., & Knottenbelt, W. (2015). The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific data*, 2(1), 1–14. <https://doi.org/10.1038/sdata.2015.7>
- [35] Australian Energy Market Operator (AEMO). (2023). *Aggregated price and demand data*. <https://aemo.com.au>
- [36] Hong, T., & Fan, S. (2016). Probabilistic electric load forecasting: A tutorial review. *International journal of forecasting*, 32(3), 914–938. <https://doi.org/10.1016/j.ijforecast.2015.11.011>
- [37] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons. https://toc.library.ethz.ch/objects/pdf/e01_978-0-470-27284-8_01.pdf
- [38] Kong, W., Dong, Z. Y., Jia, Y., Hill, D. J., Xu, Y., & Zhang, Y. (2017). Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE transactions on smart grid*, 10(1), 841–851. <https://doi.org/10.1109/TSG.2017.2753802>
- [39] Kim, T. Y., & Cho, S. B. (2019). Predicting residential energy consumption using CNN-LSTM neural networks. *Energy*, 182, 72–81. <https://doi.org/10.1016/j.energy.2019.05.230>
- [40] Oreshkin, B. N., Carpov, D., Chapados, N., & Bengio, Y. (2019). *N-BEATS: Neural basis expansion analysis for interpretable time series forecasting*. <https://doi.org/10.48550/arXiv.1905.10437>
- [41] Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and sustainable energy reviews*, 81, 1548–1568. <https://doi.org/10.1016/j.rser.2017.05.234>
- [42] Raza, M. Q., & Khosravi, A. (2015). A review on artificial intelligence based load demand forecasting techniques for smart grid and buildings. *Renewable and sustainable energy reviews*, 50, 1352–1372. <https://doi.org/10.1016/j.rser.2015.04.065>
- [43] Ribeiro, M., Grolinger, K., ElYamany, H. F., Higashino, W. A., & Capretz, M. A. M. (2018). Transfer learning with seasonal and trend adjustment for cross-building energy forecasting. *Energy and buildings*, 165, 352–363. <https://doi.org/10.1016/j.enbuild.2018.01.034>
- [44] Lu, H., Wu, J., Ruan, Y., Qian, F., Meng, H., Gao, Y., & Xu, T. (2023). A multi-source transfer learning model based on LSTM and domain adaptation for building energy prediction. *International journal of electrical power & energy systems*, 149, 109024. <https://doi.org/10.1016/j.ijepes.2023.109024>
- [45] Haben, S., Arora, S., Giasemidis, G., Voss, M., & Greetham, D. V. (2021). Review of low voltage load forecasting: Methods, applications, and recommendations. *Applied energy*, 304, 117798. <https://doi.org/10.1016/j.apenergy.2021.117798>
- [46] Khuntia, S. R., Rueda, J. L., & van Der Meijden, M. A. M. M. (2016). Forecasting the load of electrical power systems in mid-and long-term horizons: A review. *IET generation, transmission & distribution*, 10(16), 3971–3977. <https://doi.org/10.1049/iet-gtd.2016.0340>

- [47] Pinson, P. (2013). Wind energy: Forecasting challenges for its operational management. *Statistical science*, 28(4), 564–585. <https://doi.org/10.1214/13-STS445>
- [48] Koenker, R. (2017). Quantile regression: 40 years on. *Annual review of economics*, 9, 155–176. <https://doi.org/10.1146/annurev-economics-063016-103651>
- [49] Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncertainty in neural network. *International conference on machine learning* (pp. 1613-1622). Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v37/blundell15.pdf>
- [50] Gal, Y., & Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *International conference on machine learning* (pp. 1050-1059). Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v48/gal16.pdf>
- [51] Wang, Y., Zhang, N., Chen, Q., Kirschen, D. S., Li, P., & Xia, Q. (2017). Data-driven probabilistic net load forecasting with high penetration of behind-the-meter PV. *IEEE transactions on power systems*, 33(3), 3255–3264. <https://doi.org/10.1109/TPWRS.2017.2762599>
- [52] Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International journal of forecasting*, 30(4), 1030–1081. <https://doi.org/10.1016/j.ijforecast.2014.08.008>
- [53] Luo, X., Zhang, D., & Zhu, X. (2021). Deep learning based forecasting of photovoltaic power generation by incorporating domain knowledge. *Energy*, 225, 120240. <https://doi.org/10.1016/j.energy.2021.120240>
- [54] Mocanu, E., Mocanu, D. C., Nguyen, P. H., Liotta, A., Webber, M. E., Gibescu, M., & Slootweg, J. G. (2018). On-line building energy optimization using deep reinforcement learning. *IEEE transactions on smart grid*, 10(4), 3698–3708. <https://doi.org/10.1109/TSG.2018.2834219>
- [55] Kelly, J., & Knottenbelt, W. (2015). Neural nilm: Deep neural networks applied to energy disaggregation. *Proceedings of the 2nd ACM international conference on embedded systems for energy-efficient built environments* (pp. 55-64). Association for Computing Machinery (ACM). <https://doi.org/10.1145/2821650.2821672>
- [56] D’Incecco, M., Squartini, S., & Zhong, M. (2019). Transfer learning for non-intrusive load monitoring. *IEEE transactions on smart grid*, 11(2), 1419–1429. <https://doi.org/10.1109/TSG.2019.2938068>
- [57] Makonin, S., & Popowich, F. (2015). Nonintrusive load monitoring (NILM) performance evaluation: A unified approach for accuracy reporting. *Energy efficiency*, 8(4), 809–814. <https://doi.org/10.1007/s12053-014-9306-2>
- [58] Zhou, L., Li, J., Li, F., Meng, Q., Li, J., & Xu, X. (2016). Energy consumption model and energy efficiency of machine tools: A comprehensive literature review. *Journal of cleaner production*, 112, 3721–3734. <https://doi.org/10.1016/j.jclepro.2015.05.093>
- [59] Gutowski, T., Dahmus, J., & Thiriez, A. (2006). Electrical energy requirements for manufacturing processes. *13th CIRP international conference on life cycle engineering* (Vol. 31, No. 1, pp. 623-638). Leuven, Belgium. <https://www.mech.kuleuven.be/lce2006/071.pdf>
- [60] Zhang, C., Song, D., Chen, Y., Feng, X., Lumezanu, C., Cheng, W., Ni, J., Zong, B., Chen, H., & Chawla, N. V. (2019). A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, No. 01, pp. 1409-1416). Association for the Advancement of Artificial Intelligence (AAAI). <https://doi.org/10.1609/aaai.v33i01.33011409>
- [61] Zhou, T., Tang, D., Zhu, H., & Zhang, Z. (2021). Multi-agent reinforcement learning for online scheduling in smart factories. *Robotics and computer-integrated manufacturing*, 72, 102202. <https://doi.org/10.1016/j.rcim.2021.102202>
- [62] Haider, H. T., See, O. H., & Elmenreich, W. (2016). A review of residential demand response of smart grid. *Renewable and sustainable energy reviews*, 59, 166–178. <https://doi.org/10.1016/j.rser.2016.01.016>
- [63] Lago, J., De Ridder, F., & De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied energy*, 221, 386–405. <https://doi.org/10.1016/j.apenergy.2018.02.069>
- [64] Tushar, W., Yuen, C., Mohsenian-Rad, H., Saha, T., Poor, H. V., & Wood, K. L. (2018). Transforming energy networks via peer-to-peer energy trading: The potential of game-theoretic approaches. *IEEE signal processing magazine*, 35(4), 90–111. <https://doi.org/10.1109/MSP.2018.2818327>

- [65] Diagne, M., David, M., Lauret, P., Boland, J., & Schmutz, N. (2013). Review of solar irradiance forecasting methods and a proposition for small-scale insular grids. *Renewable and sustainable energy reviews*, 27, 65–76. <https://doi.org/10.1016/j.rser.2013.06.042>
- [66] Antonanzas, J., Osorio, N., Escobar, R., Urraca, R., Martinez-de-Pison, F. J., & Antonanzas-Torres, F. (2016). Review of photovoltaic power forecasting. *Solar energy*, 136, 78–111. <https://doi.org/10.1016/j.solener.2016.06.069>
- [67] Kashinath, K., Mustafa, M., Albert, A., Wu, J. L., Jiang, C., Esmaeilzadeh, S., ... & Prabhat, N. (2021). Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical transactions of the royal society a*, 379(2194), 20200093. <https://doi.org/10.1098/rsta.2020.0093>
- [68] Wang, H., Lei, Z., Zhang, X., Zhou, B., & Peng, J. (2019). A review of deep learning for renewable energy forecasting. *Energy conversion and management*, 198, 111799. <https://doi.org/10.1016/j.enconman.2019.111799>
- [69] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- [70] Xu, B., Oudalov, A., Ulbig, A., Andersson, G., & Kirschen, D. (2017). Modeling of lithium-ion battery degradation for cell life assessment. *2017 IEEE power & energy society general meeting* (pp. 1-1). IEEE. <https://doi.org/10.1109/PESGM.2017.8274005>
- [71] Severson, K. A., Attia, P. M., Jin, N., Perkins, N., Jiang, B., Yang, Z., ... & Braatz, R. D. (2019). Data-driven prediction of battery cycle life before capacity degradation. *Nature energy*, 4(5), 383–391. <https://doi.org/10.1038/s41560-019-0356-8>
- [72] Cao, J., Harrold, D., Fan, Z., Morstyn, T., Healey, D., & Li, K. (2020). Deep reinforcement learning-based energy storage arbitrage with accurate lithium-ion battery degradation model. *IEEE transactions on smart grid*, 11(5), 4513–4521. <https://doi.org/10.1109/TSG.2020.2986333>
- [73] Wan, Z., Li, H., He, H., & Prokhorov, D. (2018). Model-free real-time EV charging scheduling based on deep reinforcement learning. *IEEE transactions on smart grid*, 10(5), 5246–5257. <https://doi.org/10.1109/TSG.2018.2879572>
- [74] Milano, F., Dörfler, F., Hug, G., Hill, D. J., & Verbič, G. (2018). Foundations and challenges of low-inertia systems. *2018 power systems computation conference (PSCC)* (pp. 1-25). IEEE. <https://doi.org/10.23919/PSCC.2018.8450880>
- [75] Ajarapu, V., & Christy, C. (1992). The continuation power flow: A tool for steady state voltage stability analysis. *IEEE transactions on power systems*, 7(1), 416–423. <https://doi.org/10.1109/59.141737>
- [76] Ashton, P. M., Saunders, C. S., Taylor, G. A., Carter, A. M., & Bradley, M. E. (2014). Inertia estimation of the GB power system using synchrophasor measurements. *IEEE transactions on power systems*, 30(2), 701–709. <https://doi.org/10.1109/TPWRS.2014.2333776>
- [77] Morsi, W. G., Diduch, C. P., & Chang, L. (2010). A new islanding detection approach using wavelet packet transform for wind-based distributed generation. *The 2nd international symposium on power electronics for distributed generation systems* (pp. 495-500). IEEE. <https://doi.org/10.1109/PEDG.2010.5545860>
- [78] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning, second edition: An introduction*. MIT Press. <https://doi.org/10.1109/TPWRS.2003.821457>
- [79] Ernst, D., Glavic, M., & Wehenkel, L. (2004). Power systems stability control: Reinforcement learning framework. *IEEE transactions on power systems*, 19(1), 427–435. <https://doi.org/10.1109/TPWRS.2003.821457>
- [80] Huang, Q., Huang, R., Hao, W., Tan, J., Fan, R., & Huang, Z. (2019). Adaptive power system emergency control using deep reinforcement learning. *IEEE transactions on smart grid*, 11(2), 1171–1182. <https://doi.org/10.1109/TSG.2019.2933191>
- [81] Hafner, D., Lillicrap, T., Norouzi, M., & Ba, J. (2020). *Mastering atari with discrete world models*. <https://doi.org/10.48550/arXiv.2010.02193>
- [82] Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE transactions on systems, man, and cybernetics, part c (applications and reviews)*, 38, 156–172. <https://doi.org/10.1109/TSMCC.2007.913919>

- [83] Rashid, T., Samvelyan, M., De Witt, C. S., Farquhar, G., Foerster, J., & Whiteson, S. (2020). Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of machine learning research*, 21(178), 1–51. <http://www.jmlr.org/papers/v21/20-081.html>
- [84] Tushar, W., Saha, T. K., Yuen, C., Smith, D., & Poor, H. V. (2020). Peer-to-peer trading in electricity networks: An overview. *IEEE transactions on smart grid*, 11(4), 3185–3200. <https://doi.org/10.1109/TSG.2020.2969657>
- [85] Vázquez-Canteli, J. R., & Nagy, Z. (2019). Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied energy*, 235, 1072–1089. <https://doi.org/10.1016/j.apenergy.2018.11.002>
- [86] Garcia, J., & Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of machine learning research*, 16(1), 1437–1480. <https://www.jmlr.org/papers/volume16/garcia15a/garcia15a.pdf>
- [87] Achiam, J., Held, D., Tamar, A., & Abbeel, P. (2017). Constrained policy optimization. *International conference on machine learning* (pp. 22–31). Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v70/achiam17a>
- [88] Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., & Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 23–30). IEEE. <https://doi.org/10.1109/IROS.2017.8202133>
- [89] Berkenkamp, F., Turchetta, M., Schoellig, A., & Krause, A. (2017). Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30. https://proceedings.neurips.cc/paper_files/paper/2017/file/766ebcd59621e305170616ba3d3dac32-Paper.pdf
- [90] Molina-Markham, A., Shenoy, P., Fu, K., Cecchet, E., & Irwin, D. (2010). Private memoirs of a smart meter. *Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building* (pp. 61–66). Association for Computing Machinery. <https://doi.org/10.1145/1878431.1878446>
- [91] McKenna, E., Richardson, I., & Thomson, M. (2012). Smart meter data: Balancing consumer privacy concerns with legitimate applications. *Energy policy*, 41, 807–814. <https://doi.org/10.1016/j.enpol.2011.11.049>
- [92] Eibl, G., & Engel, D. (2014). Influence of data granularity on smart meter privacy. *IEEE transactions on smart grid*, 6(2), 930–939. <https://doi.org/10.1109/TSG.2014.2376613>
- [93] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Artificial intelligence and statistics* (pp. 1273–1282). Proceedings of Machine Learning Research (PMLR). <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [94] Taïk, A., & Cherkaoui, S. (2020). Electrical load forecasting using edge computing and federated learning. *ICC 2020-2020 IEEE international conference on communications (ICC)* (pp. 1–6). IEEE. <http://doi:10.1109/ICC40277.2020.9148937>
- [95] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM transactions on intelligent systems and technology (TIST)*, 10(2), 1–19. <https://doi.org/10.1145/3298981>
- [96] Zhang, Z., Rath, S., Xu, J., & Xiao, T. (2024). Federated learning for smart grid: A survey on applications and potential vulnerabilities. <https://doi.org/10.48550/arXiv.2409.10764>
- [97] Goodfellow, I. J., Shlens, J., & Szegedy, C. (2014). *Explaining and harnessing adversarial examples*. <https://doi.org/10.48550/arXiv.1412.6572>
- [98] Liu, Y., Ning, P., & Reiter, M. K. (2011). False data injection attacks against state estimation in electric power grids. *ACM transactions on information and system security (TISSEC)*, 14(1), 1–33. <https://doi.org/10.1145/1952982.1952995>
- [99] Chen, Y., Tan, Y., & Deka, D. (2018). Is machine learning in power systems vulnerable? *2018 IEEE international conference on communications, control, and computing technologies for smart grids (SmartGridComm)* (pp. 1–6). IEEE. <https://doi.org/10.1109/SmartGridComm.2018.8587547>
- [100] Ozay, M., Esnaola, I., Vural, F. T. Y., Kulkarni, S. R., & Poor, H. V. (2015). Machine learning methods for attack detection in the smart grid. *IEEE transactions on neural networks and learning systems*, 27(8), 1773–1786. <https://doi.org/10.1109/TNNLS.2015.2404803>

- [101] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proceedings of the IEEE*, 107(8), 1738–1762.
<https://doi.org/10.1109/JPROC.2019.2918951>
- [102] Han, S., Mao, H., & Dally, W. J. (2015). *Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding*. <https://doi.org/10.48550/arXiv.1510.00149>
- [103] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877–1901.
https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- [104] Bubeck, S., Chadrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., ... & Zhang, Y. (2023). *Sparks of artificial general intelligence: Early experiments with gpt-4*. https://scalar.usc.edu/works/simulated-worlds/media/Sparks_of_AGI_and_What_are_World_Models.pdf
- [105] Das, A., Kong, W., Sen, R., & Zhou, Y. (2023). *A decoder-only foundation model for time-series forecasting*. <https://doi.org/10.48550/arXiv.2310.10688>
- [106] Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE access*, 8, 108952–108971. <http://doi.org/10.1109/ACCESS.2020.2998358>
- [107] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2018). Digital twin in industry: State-of-the-art. *IEEE transactions on industrial informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
- [108] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems* (Vol. 27, pp. 2672–2680). Curran Associates Inc.
https://proceedings.neurips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html
- [109] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [110] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
<https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>
- [111] Jain, S., & Wallace, B. C. (2019). Attention is not explanation. *Proceedings of the 2019 conference of the north american chapter of the association for computational linguistics: Human language technologies, volume 1 (long and short papers)* (pp. 3543–3556). Minneapolis, Minnesota: Association for Computational Linguistics.
<https://doi.org/10.18653/v1/N19-1357>
- [112] European Commission. (2021). *Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts* (COM/2021/206 final). EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- [113] Kostic, T., Preiss, O., & Frei, C. (2005). Understanding and using the IEC 61850: A case for meta-modelling. *Computer standards & interfaces*, 27(6), 679–695. <https://doi.org/10.1016/j.csi.2004.09.008>
- [114] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in neural information processing systems*, 28.
https://proceedings.neurips.cc/paper_files/paper/2015/file/86df7dcfd896fcdf2674f757a2463eba-Paper.pdf
- [115] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2), 182–197.
<https://doi.org/10.1109/4235.996017>
- [116] Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. *IEEE transactions on knowledge and data engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>

