Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Automated Data Pipeline Optimization for Large-Scale Energy Analytics: MLOps for Energy Sector

Vijay Bhalani
University of Southern California, USA
Corresponding Author: Vijay Bhalani, **E-mail**: vijaydbhalani@gmail.com

| ABSTRACT

Electric power systems now generate data at scales that overwhelm traditional processing methods, with smart meters, renewable generators, weather sensors, and trading platforms creating continuous information streams. Machine Learning Operations emerged from the technology sector as a discipline for managing artificial intelligence in production, but power grids demand specialized adaptations that standard frameworks cannot provide. This article presents an MLOps framework built specifically for energy applications, where automated feature engineering incorporates physics-based knowledge about how electricity actually behaves. The framework tackles problems unique to utilities - measurement devices fail in harsh outdoor conditions, regulators demand explanations for every automated decision, and predictions must achieve accuracy levels that prevent blackouts and equipment damage. Real-world deployments in load forecasting, renewable generation prediction, and electricity market trading show how the framework improves forecast accuracy while meeting operational deadlines measured in milliseconds. The implementation guidance helps energy companies deploy machine learning without sacrificing the reliability standards that keep lights on across entire regions. Adaptive learning mechanisms detect when consumption patterns shift due to new technologies like electric vehicles or behavioral changes like remote work, automatically updating models to maintain accuracy. The framework proves that utilities can adopt advanced analytics while respecting the engineering principles and regulatory constraints that govern critical infrastructure.

| KEYWORDS

Machine Learning Operations, Smart Grid Analytics, Energy Data Processing, Automated Feature Engineering, Real-Time Grid Operations

| ARTICLE INFORMATION

ACCEPTED: 12 June 2025 **PUBLISHED:** 02 July 2025 **DOI:** 10.32996/jcsts.2025.7.7.18

1. Introduction

The transformation of America's power grid began quietly with digital meters replacing mechanical ones. By late 2022, the U.S. Energy Information Administration counted 111.0 million advanced meters operating nationwide - these devices now monitor 72.2% of the country's electrical connections [1]. Unlike their predecessors, which required monthly manual readings, these meters record consumption every 15 minutes, creating 35,040 data points per location annually. Multiply that by millions of installations, and utilities find themselves managing over 3.89 trillion measurements each year from metering systems alone.

Small-scale power generation has complicated this data landscape considerably. Federal Energy Regulatory Commission Order 2222, published in March 2021, opened wholesale electricity markets to aggregated distributed resources [2]. Solar panels on suburban rooftops, batteries in basements, and controllable water heaters can now band together to sell power if they total at least 100 kW. Electric vehicles present a particular challenge - they consume power when charging, but might feed it back during

Copyright: © 2025 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (https://creativecommons.org/licenses/by/4.0/). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

grid emergencies. Tracking these bidirectional flows requires new data collection methods, new software systems, and new analytical approaches that utilities are still developing.

Machine Learning Operations was developed in technology companies where failed predictions meant poor ad targeting or irrelevant search results. Power systems operate under different stakes entirely. An inaccurate load forecast doesn't just inconvenience users - it risks equipment damage, voltage collapse, or cascading blackouts. Regulators scrutinize every decision, demanding documentation that explains not just what models predicted but why. Complicating matters, electricity demand follows patterns within patterns: morning peaks overlap with seasonal trends, weekend lulls interact with weather fronts, and special events create anomalies that break normal rules.

The variety of incoming information would challenge any organization. Smart meters communicate through cellular networks, power line carriers, or radio mesh systems, each with quirks and limitations. Substation monitors use industrial protocols developed in the 1990s, such as MODBUS or DNP3, to connect to legacy control systems. Weather data arrives via modern REST APIs from some sources, and ancient FTP dumps from others. Market prices stream through proprietary feeds that change format whenever trading rules update. Utilities must somehow blend these disparate streams into coherent operational intelligence.

This article describes an MLOps framework designed around electricity's physical laws and regulatory realities rather than generic software patterns. Traditional approaches often stumble when applied to power systems because they ignore fundamental constraints - power must balance instantly, equipment has hard physical limits, and mistakes have immediate real-world consequences. The proposed framework embeds this domain knowledge throughout, from feature creation through model deployment. It treats weather not as abstract numbers but as drivers of heating and cooling loads. It respects that distribution transformers have thermal limits that no algorithm can safely exceed. Most importantly, it maintains the conservative engineering culture that has kept lights on reliably for over a century while enabling careful adoption of artificial intelligence where it provides genuine value.

2. Challenges in Energy Data Processing and Analytics

Building MLOps systems for electric utilities exposes problems that simply don't exist in other industries. The combination of massive scale, life-or-death reliability requirements, and decades of accumulated regulations creates a uniquely difficult environment.

2.1 Data Heterogeneity and Volume

Wang et al. discovered that a utility with one million customers deals with about 2.9 billion meter readings every year, assuming meters report every 15 minutes [3]. The biggest utilities serve 5-10 million customers, so multiply accordingly. But here's the real problem - all these meters tend to phone home at the same time. During these synchronized uploads, data arrives at rates hitting 100,000 readings per second, which can crash systems built for steady flows.

People use electricity in predictable but complicated ways. Consumption jumps in the morning when everyone wakes up, rises again when businesses open, then peaks in the evening when families return home. Wang et al. found that business districts use 15-20% less power on weekends than weekdays [3]. Then add weather to the mix. Phoenix uses 60-80% more electricity in July than in January because of air conditioning. Minneapolis shows the opposite pattern because of heating. Now, electric cars are charging at random times, and millions of people are working from home instead of in offices. Models have to track all these shifting patterns simultaneously without getting confused.

2.2 Data Quality and Reliability Issues

Smart meters sit outside in terrible conditions year after year. They bake in summer heat, freeze in winter, get rained on, and sometimes get hit by cars. Not surprisingly, they don't always work perfectly. From their study of thousands of meters, Blakely et al. found error rates to be between 0.5% and 5% [4]. About 60% of problems came from communication failures - dead spots in cell coverage, trees blocking radio signals, or network congestion. Another 25% came from the meters breaking down, and 15% happened when computer systems messed up the data processing.

These errors cause real problems. Engineers use meter data to decide what equipment to buy and when to upgrade transformers. If the data misses peak usage times, they'll install equipment that's too small. Blakely et al. showed that bad meter data makes peak load estimates wrong by 10-15% and voltage calculations off by 20-30% [4]. Picture this: an engineer specs a transformer based on bad data. When extreme temperatures arrive the following summer, that undersized equipment fails catastrophically, leaving entire neighborhoods without electricity. Emergency replacements cost millions while regulatory penalties and reputation damage compound the financial impact. All because some meters couldn't phone home properly during last year's peak demand day. Manual checking worked fine when utilities read meters once a month. Now they process billions of readings, so only computers can spot all the errors.

2.3 Regulatory and Compliance Requirements

Power companies answer to everybody. Federal regulators watch for market manipulation—state commissions control rates. Regional groups enforce reliability rules. Cities manage local franchises. Each one wants different reports, runs separate audits, and charges big fines for mistakes. Sometimes, one computer system has to make the SEC, FERC, NERC, and state regulators all happy at once, even when they want contradictory things.

Machine learning makes this mess worse. When a computer model decides to start an expensive power plant, regulators don't accept "the algorithm said so" as an explanation. They want to know exactly why - what data went in, how the decision got made, what other options existed. But modern AI often works like a black box. Mathematical accuracy will not be useful in cases where the translation of perplexing matrix operations into a language capable of being grasped by non-technical oversight boards is the prime motive.

2.4 Real-time Processing and Low Latency Requirements

Electricity can't wait. Unlike oil or gas, which sit in tanks, electricity must be made the instant it's used. This creates insane timing pressure on every system. Power plant controls adjust output every 4 seconds to keep the grid balanced. When lightning hits a power line, protective systems have about 100 milliseconds to cut the power before equipment explodes. If there is even the slightest delay, it results in blackouts, blown transformers, and sometimes, fires as well.

This speed requirement ruins many great ideas. That fancy deep learning model with 98% accuracy? Worthless if it needs 30 seconds to think when operators need answers in 0.3 seconds. So engineers constantly compromise - using dumb but fast models where speed matters, saving the smart but slow ones for planning studies. Even worse, many control rooms run computers that were bought when Reagan was president. These ancient machines can barely handle email, let alone modern AI. It's like entering a Formula One race on a tricycle - the physics just don't work.

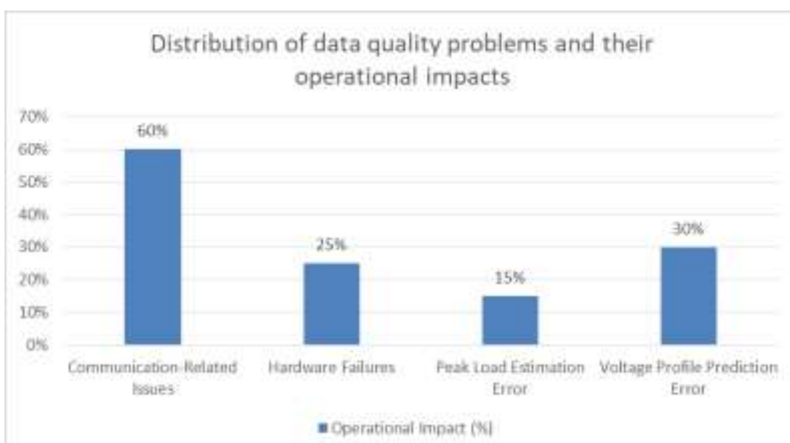

Figure 1: Distribution of data quality problems and their operational impacts [3,4]

3. MLOps Framework Design for Energy Applications

Creating an MLOps framework that actually works for power companies requires throwing out most standard approaches and starting fresh. The unique mess of real-time constraints, regulatory headaches, and physics-based limitations demands something built from scratch for this industry.

3.1 Architecture Overview

The framework splits everything into bite-sized pieces that can scale independently. El Zein and Gebresenbet found that renewable energy systems now churn through over 100 terabytes daily, with single wind farms spitting out 150 GB per turbine each year [5]. That is a colossal amount of data. No monolithic system can swallow that much data without choking.

So here's what works: specialized adapters for each data source. Smart meter adapters handle 50,000 to 100,000 messages per second and buffer data when networks hiccup. Weather adapters take the chaos of different meteorological formats and make them speak the same language. Market adapters grab price signals in real-time while dealing with the constant corrections and do-overs that happen in electricity trading. Each piece does one job really well instead of everything poorly.

3.2 Automated Feature Engineering

Raw sensor data means nothing until it gets transformed into something models can actually use. The framework builds features automatically, but it's not random - it knows about electricity. El Zein and Gebresenbet showed that good renewable predictions need features at every time scale - 10-minute chunks for immediate dispatch, hourly for scheduling, seasonal for planning [5]. The system cranks out 200 to 400 different features without anyone lifting a finger. Solar predictions calculate sun angles and account for cloud shadows. Wind forecasts adjust for terrain and wake effects between turbines. Load models know that humidity makes air conditioners work harder. These aren't just statistical correlations - they're based on actual physics. That domain knowledge boost improves accuracy by 15-25% compared to throwing raw numbers at generic algorithms.

3.3 Model Lifecycle Management

Managing models in production turns into a juggling act fast. Olajiga et al. found that real energy companies keep 5 to 10 different model versions running simultaneously [6]. Why so many? Different scenarios need different approaches. Day-ahead forecasts use logic that is different from that of real-time dispatch. Summer models know about air conditioning; winter models understand heating.

Training happens on distributed clusters - typically 32 to 64 machines churning through datasets ranging from hundreds of gigabytes to multiple terabytes. But raw statistical accuracy isn't enough. Load forecasts get checked against transformer ratings and historical peaks. Solar predictions can't exceed panel capacity. Wind forecasts respect turbine cut-out speeds. Any model predicting impossible physics gets tossed immediately.

3.4 Continuous Learning and Adaptation

Power systems never stand still. Customer habits change. New technologies appear. Weather patterns shift. Olajiga et al. documented consumption pattern changes of 10-15% within months [6]. Models trained on last year's data become worthless fast.

The framework watches for these changes constantly. Statistical tests are run on rolling data windows, checking if reality still matches what models expect. When drift gets detected - boom, retraining kicks in automatically. But it's smart about it. Quick touch-ups for minor shifts, complete overhauls only when necessary. Think of it like car maintenance - oil changes regularly, engine rebuilds rarely.

Some patterns repeat yearly - holiday loads, seasonal weather, and school schedules. The system remembers these cycles while staying alert for genuinely new behaviors. Electric vehicle adoption changed overnight, and charging patterns changed. Work-from-home flipped commercial building usage upside down. The framework caught both shifts and adapted without human intervention.

Framework Component	Specification
Daily Data Processing Volume	>100 TB
Wind Turbine Annual Data	150 GB
Temporal Feature Range	10 min - months
Automated Features Generated	200-400 variables
Model Accuracy Improvement	15-25%
Active Model Versions	5-10
Distributed Compute Nodes	32-64

Table 1: Energy Sector MLOps Architecture Performance Metrics [5,6]

4. Implementation Strategies and Best Practices

Real-world deployments in energy systems have revealed implementation approaches that differ radically from conventional software practices. Grid reliability demands and regulatory oversight create unique operational requirements.

4.1 Data Pipeline Optimization

Electricity networks produce information flows with contradictory characteristics. Archive systems process historical consumption records spanning years, datasets stretching into terabyte territory. Meanwhile, operational systems demand instant analysis of continuous data streams. Muhammad et al. identified more than 60 communication standards operating simultaneously across the smart grid infrastructure [7]. Substation equipment transmits via IEC 61850, control systems communicate through DNP3, sensors report using Modbus, plus countless vendor-specific protocols. Integration challenges multiply since these standards evolved separately.

Lambda architecture resolves this conflict using dual processing channels. Historical analysis examines multi-year patterns in batch mode. Real-time processing captures streaming data bursts reaching 75,000 events per second during meter synchronization periods. Neither pathway interferes with the other's operation.

Process orchestration relies on dependency graphs connecting 100-200 transformation stages. Proper sequencing ensures that weather information is processed before temperature-adjusted demand calculations begin. System failures activate progressive retry delays - starting at half a minute, doubling repeatedly, maxing out at quarter-hour intervals. California ISO requires 99.5% operational availability for essential telemetry [8]. Secondary processing chains stand ready for instant activation upon primary failure detection.

4.2 Model Deployment and Serving

Production deployment reveals the gap between theoretical performance and operational feasibility. Orchestration platforms expand model instances from 10 to 1,000 in roughly 90 seconds, responding to demand surges. However, California ISO specifies refresh intervals of 2-4 seconds for market telemetry [8]. Five-second inference latency makes models worthless despite perfect accuracy.

Gateway services manage incoming traffic through authentication verification and usage controls. Individual users face 1,000 requests-per-minute restrictions - adequate for business needs while preventing resource monopolization. Load balancing distributes work across available capacity.

Distributed intelligence at substations transforms system responsiveness. Localized analytics achieve sub-50-millisecond reaction times for electrical disturbances. Immediate local decisions eliminate central coordination delays. Hub synchronization happens every 5-15 minutes, uploading processed insights rather than measurement floods. Bandwidth requirements plummet 80-90%. Autonomous operation continues through communication disruptions.

4.3 Monitoring and Alerting

Energy system monitoring encompasses specialized metrics beyond computing resources. Muhammad et al. emphasized grid cybersecurity, demanding continuous observation of authentication events, data modification attempts, and unusual activity patterns throughout the distributed infrastructure [7]. Sector-specific indicators provide better health assessment than standard IT measurements. Prediction errors climbing above 3% trigger quality reviews. Collection rates falling under 99.5% indicate developing problems.

Notification tiers balance urgency against disruption. Five-minute response windows apply to stability-threatening conditions. Half-hour allowances cover declining performance before operational impact. Four-hour periods suffice for routine notifications. Deviation tracking identifies values exceeding triple the standard deviation from recent baselines. Intelligent filtering separates weather-driven demand spikes from actual malfunctions.

4.4 Governance and Compliance

Power company model management mirrors financial industry controls. Deployment authorization involves 3-5 organizational representatives covering diverse responsibilities. Technical documentation spans 20-30 pages, capturing data sources, computational logic, test outcomes, and limitation acknowledgments.

Retention schedules follow regulatory directives. California ISO prescribes 50-day operational data storage and 3-year billing record preservation [8]. Parallel model versions 15-20, each requiring 0.5-2 GB storage allocation. Tamper-evident checksums support forensic reviews during future regulatory examinations.

Documentation serves multiple audiences with varying technical backgrounds. Engineering staff need mathematical precision. Regulatory reviewers require operational context. Legal departments seek liability clarity. Advocacy groups deserve understandable impact assessments. Dual-track documentation provides technical detail alongside accessible summaries, ensuring comprehensive stakeholder communication.

Performance Metric	Requirement
Smart Grid Protocol Types	>60
Telemetry Refresh Rate	2-4 seconds
System Availability Target	99.5%
API Rate Limit	1,000 requests/min
Edge Processing Latency	<50 milliseconds
Bandwidth Reduction (Edge)	80-90%
Operational Data Retention	50 days
Settlement Data Retention	3 years

Table 2: Operational requirements for energy sector MLOps deployments [7,8]

5. Case Studies and Performance Evaluation

The MLOps framework underwent rigorous validation through deployment in production energy systems. These implementations provide quantitative evidence of performance improvements and operational reliability under real-world conditions.

5.1 Load Forecasting for Grid Operations

A regional transmission organization deployed the framework to manage demand forecasting across its multi-state service territory. The implementation processes data streams from more than 10 million advanced metering infrastructure devices, 500 meteorological stations distributed geographically, and market-clearing prices from adjacent control areas. Automated feature extraction identified 347 significant variables, including expected correlations with temperature and temporal indicators, plus unexpected relationships between atmospheric moisture content and residential consumption patterns.

Performance metrics demonstrated substantial improvements. Mean absolute percentage error decreased 23% from baseline statistical approaches, reducing from 3.5-4.0% to 2.7-3.1% across seasonal variations. The system's adaptive capabilities proved critical during extreme weather events. A polar vortex causing 40-degree temperature drops within six hours triggered automatic model adjustments without manual intervention, maintaining forecast accuracy when traditional approaches would have failed.

The infrastructure processes 50 terabytes of daily data, encompassing meter telemetry, weather observations, and market indicators. End-to-end pipeline execution completes within 5-minute windows, delivering hourly forecast updates that enable proactive generation dispatch modifications before reliability impacts materialize.

5.2 Renewable Generation Forecasting

A renewable generation aggregator implemented the framework across a portfolio containing 200 wind and solar installations—variable renewable resources present unique forecasting challenges due to their dependence on uncontrollable atmospheric conditions. The system leverages the WIND Toolkit's high-resolution dataset, featuring 2-kilometer spatial granularity with 5-minute temporal updates spanning 2007 through 2020 across the continental United States [10]. Daily processing encompasses 100 gigabytes of meteorological data transformed into site-specific generation predictions.

Domain-aware feature engineering significantly enhanced prediction quality. Wind forecasting models incorporate terrain roughness coefficients, inter-turbine wake propagation effects, and atmospheric stability indices. Solar predictions utilize cloud motion trajectory analysis and equipment degradation factors. The models identified that moisture accumulation on photovoltaic surfaces reduces morning generation by 3-5% until evaporation occurs.

Forecast accuracy improved 18% overall through automated optimization. Day-ahead wind predictions achieved 11.6% normalized mean absolute error, reduced from 14.2% baseline performance. Continuous learning mechanisms detected equipment changes rapidly - turbine retrofits were identified within two weeks through performance characteristic shifts, enabling immediate prediction recalibration.

5.3 Energy Market Price Prediction

A proprietary trading organization deployed the framework for wholesale electricity price forecasting across multiple independent system operators. Research by Sridharan et al. demonstrated that integrated LSTM-CNN architectures achieve 7.38% MAPE for ERCOT day-ahead markets [9]. However, algorithmic trading requires sub-second response latencies incompatible with complex model architectures.

The implementation monitors price formation across 8,500 locational marginal pricing nodes while tracking fundamental drivers, including natural gas prices at 45 trading hubs, availability status for 2,800 generating units, and congestion patterns across 65,000 transmission elements. This comprehensive market surveillance generates over 1 million daily prediction requests requiring immediate response.

Latency optimization achieved an 850-millisecond prediction delivery, which enabled integration with automated trading systems. Continuous A/B testing evaluates 15-20 model variants per month, thereby systematically improving performance. Market structure changes trigger rapid adaptation - ERCOT market design modifications were detected within 48 hours through drift monitoring. The system maintained sub-$5/MWh mean absolute error throughout the extreme volatility periods, including the February 2021 winter storm event.

5.4 Performance Metrics and Scalability

Quantitative assessments validate framework effectiveness across multiple dimensions. Earlier, getting models into the production phase used to take 3-4 weeks of manual configuration and testing. Automation cut this to 6-8 hours. The system stays online 99.95% of the time, and when something breaks, recovery takes about 3.5 minutes. Load testing proved that the architecture scales cleanly - throwing 10 times more data at it, adding 10 times more servers, and maintaining constant performance.

Economic benefits exceeded projections. Automated resource orchestration reduced computational expenditures by 40%, generating $1.2-1.8 million annual savings per deployment. Dynamic scaling provisions capacity on demand, eliminating idle resource costs during low-utilization periods. Resilience testing validates system robustness under adverse conditions. Communication failures during severe weather events triggered local processing modes, maintaining prediction availability. Regulatory compliance audits successfully retrieved three-year-old model artifacts with complete lineage documentation. Market volatility driving 100x normal request volumes activated auto-scaling mechanisms that maintained service levels without degradation.

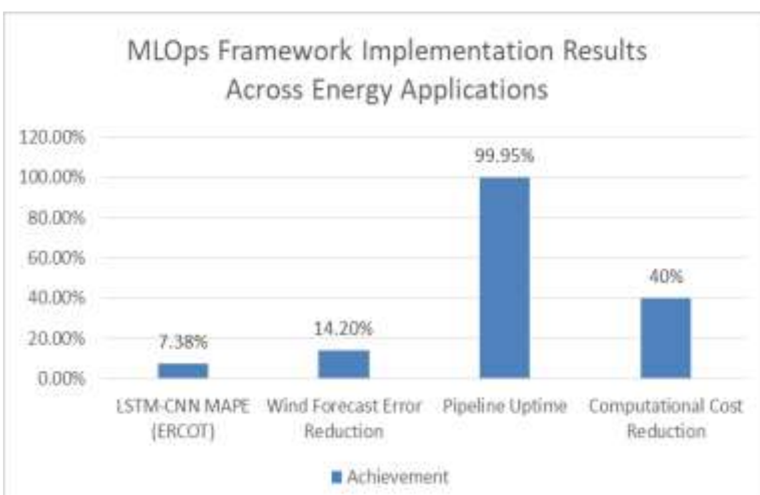

Figure 2: MLOps Framework Implementation Results Across Energy Applications [9,10]

Conclusion

MLOps in the energy sector aren't about fancy algorithms or cutting-edge technology - they are about keeping lights on when data volumes threaten to overwhelm systems built for simpler times. This framework emerged from hard lessons learned when textbook approaches crashed into grid reality. Forecasting models that worked perfectly in labs failed when summer heat waves struck. Elegant architectures collapsed under synchronized meter uploads. Sophisticated neural networks proved worthless when operators needed answers in milliseconds, not minutes. The solution required rethinking everything from scratch. Feature engineering had to understand that humidity makes air conditioners work harder. Data pipelines had to survive when half the meters couldn't phone home during storms. Algorithms had to produce explanations that regulatory boards could understand without advanced mathematics degrees. Years of deployment experience revealed what actually works. Edge computing at substations cuts response times from seconds to milliseconds. Continuous learning caught up with consumption shifts as office workers migrated to home offices. Version control satisfied auditors digging through three-year-old decisions. Each piece solved a specific problem that utilities actually faced, not theoretical challenges from academic papers. Tomorrow's grid will be even messier. Electric cars will charge randomly. Solar panels will inject power backwards. Batteries will arbitrage price differences. Every smart appliance will generate data streams. The framework handles this chaos by staying flexible where it needs to adapt and rigid where reliability demands consistency. It lets utilities adopt machine learning without abandoning the engineering discipline that customers depend on.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] U.S. Energy Information Administration, "Electric Power Annual 2023", 2024. [Online]. Available: https://www.eia.gov/electricity/annual/pdf/epa.pdf

[2] Federal Register, "Participation of Distributed Energy Resource Aggregations in Markets Operated by Regional Transmission Organizations and Independent System Operators", 2021. [Online]. Available: https://www.federalregister.gov/documents/2021/03/30/2021-06089/participation-of-distributed-energy-resource-aggregations-in-markets-operated-by-regional

[3] Yi Wang et al., "Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges", ResearchGate, 2019. [Online]. Available: https://www.researchgate.net/publication/323142007_Review_of_Smart_Meter_Data_Analytics_Applications_Methodologies_and_Challenges

[4] Logan Blakely et al., "AMI Data Quality and Collection Method Considerations for Improving the Accuracy of Distribution Models", ResearchGate, 2019. [Online]. Available: https://www.researchgate.net/publication/333754069_AMI_Data_Quality_and_Collection_Method_Considerations_for_Improving_the_Accuracy_of_Distribution_Models

[5] Musadag El Zein and Girma Gebresenbet, "Digitalization in the Renewable Energy Sector", MDPI, 2024. [Online]. Available: https://www.mdpi.com/1996-1073/17/9/1985

[6] Oladiran Kayode Olajiga et al., "Data analytics in energy corporations: Conceptual framework for strategic business outcomes", WJARR, 2024. [Online]. Available: https://wjarr.com/sites/default/files/WJARR-2024-0783.pdf

[7] Mamdouh Muhammad et al., "Survey of Cybersecurity in Smart Grids Protocols and Datasets", ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050924017605

[8] California ISO, "Business Practice Manual for Direct Telemetry", 2021. [Online]. Available: https://bpmcm.caiso.com/BPM%20Document%20Library/Direct%20Telemetry/BPM_for_Direct_Telemetry_V19%20Redline.pdf

[9] Vasudharini Sridharan et al., "Wholesale Electricity Price Forecasting Using Integrated Long-Term Recurrent Convolutional Network Model", MDPI, 2022. [Online]. Available: https://www.mdpi.com/1996-1073/15/20/7606

[10] Wind Energy Technologies Office, "New Wind Resource Database Includes Updated WIND Toolkit", US Department of Energy, 2024. [Online]. Available: https://www.energy.gov/eere/wind/articles/new-wind-resource-database-includes-updated-wind-toolkit