

A Thematic Review of AI-Based Energy Forecasting Approaches: Accuracy, Uncertainty, and Economic Value

Kunshraj Singhania

Abstract: *Power grids everywhere now rely more heavily on renewables, turning forecasts once used just for planning into daily necessities. Not like coal or gas plants, wind and sunlight shift constantly, shaped by weather patterns and daylight changes. Because output jumps around unpredictably, keeping supply steady demands extra reserves - costs climb along with uncertainty. To handle swings better, utilities lean increasingly on smart computer methods: neural networks, ensemble tricks, statistical hybrids, pattern-spotting algorithms - all trained to guess sun strength, wind speed, how much people will use, even price shifts. Yet here's what studies rarely stress - the gap between lab results boasting higher scores and real-world reports talking about messy databases, unclear ownership rules, shaky profit numbers when firms actually roll out these tools. Looking at studies from 2020 to 2025, this review brings together academic and non-academic work exploring if AI-powered energy forecasts beat older statistical methods in cost terms. Focus lands first on predicting renewable output, though insights about demand and pricing pop up too along the way. Instead of old-school stats, many now turn to machine learning - sometimes shallow, sometimes stacked deep. Data cleanliness matters more than clever code, often deciding success before any model runs. Picking between simpler algorithms and complex neural nets involves balancing effort, transparency, and results case by case. Guessing ranges beats single numbers when dealing with sun and wind, so measuring doubt has become standard practice. Better predictions only count if they save money or cut waste, which ties progress back to real-world gains eventually. Nowhere is it guaranteed that tighter error scores mean better real-world results. Even when machines beat older techniques on measures like Mean Absolute Error (Benti et al., 2023; Abisoye, 2024) - how far off predictions are on average - the edge can vanish in practice. Root Mean Square Error, which stresses bigger mistakes more, often favors AI, yet this does not automatically lead to savings or efficiency. Mean Absolute Percentage Error shows percent-level gaps, a common yardstick, but high marks here mislead if deployment costs rise too much. What plays out in spreadsheets may fail under live conditions. Gains shrink when forecasts stretch further ahead in time. Systems able to adjust quickly benefit more from precise inputs. Rules set by regulators shape how useful any prediction method becomes. Heavy computing needs eat into advantages unless hardware keeps pace. Older grids with weak data networks gain little, even with advanced tools running behind the scenes. Looking at India offers clues because of how fast it adds renewables (IEA, 2021; NREL, n.d.) alongside rules tying forecast errors to money consequences. Though machines predict better, they save cash just when systems align a certain way. Coming studies need to weigh costs in comparisons, speed up number crunching, also blend uncertainty estimates into daily choices more tightly.*

Keywords: Artificial Intelligence (AI), Energy Forecasting, Renewable Energy Generation, Machine Learning, Deep Learning, Probabilistic Forecasting, Smart Grids

1. Introduction

Most of the time, electricity gets used right when it's made, because storing big amounts is tough and costs a lot. So grids rely on steady adjustments throughout each day. Matching output to usage happens nonstop to prevent disruptions. Predicting what comes next shapes how power flows minute by minute. Looking at past patterns plus weather helps shape those guesses. Forecasting isn't just number crunching - it drives actual decisions. What happened yesterday gives clues about today's needs. These predictions keep lights on without delay.

Long ago, guesses about power needs focused on how much electricity people would use. That kind of guess goes by another name: load forecasting - figuring out how much will be pulled from outlets in hours or days ahead. Power came mostly from steady sources like coal, gas, and water-driven turbines. Those old setups ran without many surprises. Now things shift fast because sunlight panels and spinning wind towers grow everywhere. This twist changes everything about what numbers come next. Sun and wind do not push electrons at a constant rate. Their strength dances with weather patterns. Brightness levels matter. So does how hard air moves through blades. Clouds drift. Heat rises. All shape the flow. These shifts earn a label: Variable Renewable Energy. What comes

out depends on nature's rhythm. Predictions now chase moving targets.

Most days, things run fine - until they do not. If solar or wind produce less than expected, backup sources must step in, often at high cost. A surprise shortfall means last-minute deals, higher prices. On the flip side, too much clean energy can flood the system, forcing shutdowns even when machines could keep running. Cutting production on purpose wastes money, simple as that. Mistakes in guesses become bills, paid by someone, somehow.

More complicated energy setups now use smart computer methods to predict outcomes. Think of these tools as programs that learn from information without being told every step. One kind uses decision trees built together - called Random Forest - or ones boosted through repeated refinement like Gradient Boosting. Another type relies on brain-inspired layers, stacking them deep; LSTM is an example here. They find twisted links among factors, where changes aren't steady or clear-cut. Old math tricks often fail when connections twist too much.

Most studies show AI forecasts beat older math-based ones like ARIMA (Benti et al., 2023; Krishnamurthy, 2024), which rely only on past numbers and patterns. Deep networks, it turns out, tend to score better than basic stats (Abisoye, 2024;

Volume 15 Issue 8, August 2026

Fully Refereed | Open Access | Double Blind Peer Reviewed Journal

www.ijsr.net

Yandar et al., 2024) in predicting solar or wind output - judged by common measures such as MAE, RMSE, and MAPE, according to newer papers in energy journals. Still, more accurate predictions don't always mean bigger gains once real-world costs enter the picture.

One persistent theme across research involves what some call the accuracy-economics divide. While scholarly work often focuses on improving prediction quality, showcasing slight gains in error rates, practical applications tell a different story. Deployment, it turns out, depends heavily on resources beyond algorithms. Infrastructure demands emerge clearly in real-world settings - cloud platforms, secure data pipelines, skilled personnel. Such requirements appear consistently in sector analyses (McKinsey & Company, 2024; Deloitte, 2023). Expenses tied to these elements tend to grow quickly.

Reports from consultants and trade groups tend to present artificial intelligence through the lens of wide-scale digital change. A 2024 McKinsey analysis highlights utility applications where prediction tools, efficiency models, and equipment insights work together, aiming at stronger operations. In parallel, Deloitte's 2023 review underlines how linked data environments and unified platforms help unlock benefits within power networks. Such views imply gains in predictive accuracy matter most once woven into daily processes.

This difference brings up a key issue: Does artificial intelligence predict more cost-effectively than standard techniques, yet deliver only better numbers?

To address this question, shifting focus away from metrics allows a clearer view of how better forecasts affect finances. Value emerges not solely through one path (DeepMind, 2019; IEA, 2021) - multiple routes exist where accuracy impacts results

- Lower fines for imbalances appear where rules apply. Where oversight exists, financial corrections shrink.
- Lower reserve procurement costs.
- Trading choices become more effective through better analysis. Hedging strategies adapt with clearer market signals.
- Reduced renewable curtailment.
- Greater stability in power supply reduces financial losses from blackouts.

Still, how large those gains are relies heavily on how markets are structured, alongside regulations that shape them. Renewable energy's share plays a role, tied closely to the grid's ability to adapt. Rules set by authorities matter just as much as the technology involved. System responsiveness becomes critical when supply fluctuates. Market frameworks define what outcomes can emerge under stress.

Forecasting energy demand splits mainly by how far ahead predictions stretch. Minutes to hours out? That kind of near-instant outlook guides immediate power dispatch. Looking one day forward helps shape bids in electricity markets, also shapes which generators get switched on. Further ahead, weeks or years, the projections shift toward infrastructure plans, funding choices. Time scale changes everything - what matters now fades later; what seems distant gains weight.

Because they handle patterns across time steps well, deep learning tools like LSTM - recurrent networks built for sequences - are often chosen for near-future predictions. Recently introduced into this space, Transformers, which rely on attention mechanisms and started in language tasks, now serve forecasting by capturing distant connections without struggling over time gaps.

Though deeper models often need more data, they also demand stronger computing capacity than traditional approaches like Random Forest - a method merging many decision trees - or Support Vector Machines, which locate best-fit separation lines under supervision. What we mean by computational complexity is how much processor strength and memory space a model uses during training and deployment. When complexity rises, so do running expenses along with environmental impact, particularly if updates happen often.

Still, more complex models often face greater hurdles during real-world implementation.

Despite its scale, India's power system faces unique challenges due to fast-growing renewable inputs. Solar and wind installations have risen sharply over recent years. Because generation fluctuates, balancing supply becomes complex. Instead of relying on traditional methods, the nation introduced REMCs (NSGM, 2022; NSGM, 2017) - specialized hubs that merge predictive models with live data flow across networks. These centres support stability by anticipating shifts before they disrupt service. Strategic positioning makes India a meaningful example when assessing cost-driven arguments around forecast systems. Official records from national legislative sources confirm their integration into daily operation.

Financial consequences follow forecasting mistakes within India's Deviation Settlement Mechanism structure (IEA, 2021; Prayas Energy Group, n.d.). When predictions fail, payment obligations emerge under current rules. Accuracy now weighs heavier on finances due to evolving regulations around renewable output reporting. Stricter standards appear likely ahead, shaping how forecasts influence revenue outcomes - per Reuters in 2025.

This setting allows observation of AI forecasting through real-world conditions: as prediction quality advances, corresponding shifts in financial outcomes become visible. Still, a constraint emerges - not all gains translate equally, since atmospheric unpredictability restricts how precise forecasts can ever be. When accuracy nears that limit, further enhancements contribute less to overall value.

This thematic review aims to compile recent scholarly reviews and grey literature on AI based energy forecasting to evaluate whether such methods are economically more suitable when compared to traditional statistical methods. The review will focus primarily on renewable energy generation forecasting while also including load and price forecasting for deeper comparison.

The paper is organized around recurring themes identified across sources

- Change from statistical to AI based methods
- The role of feature engineering
- Trade-offs between mechanical learning and deep learning
- Probabilistic forecasting and uncertainty representation.
- Translation of forecasting accuracy into economic value.

By comparing academic findings with industry narratives, this review aims to clarify when AI forecasting delivers meaningful economic advantage and when its benefits may be overstated.

2. Literature Review

Forecasting Taxonomy and Baseline Approaches

Across varying timespans, market setups, and functional aims, energy prediction involves linked challenges rather than one isolated task. Generation estimates, demand projections, and forecasts for electricity costs form the trio seen most often. While differing in data needs and target outcomes, these types aim at the same result - less unpredictability in how grids run and develop. What ties them together isn't method but purpose: shaping decisions with clearer future signals.

Generation Forecasting

Forecasting what comes next begins with estimating how much electricity power stations will produce, particularly those using renewables like wind or sunlight. Since sun- and wind-based production shifts with atmospheric patterns, predictions typically pair past performance records alongside climate factors including:

- Solar Irradiance,
- Wind Speed,
- Humidity,
- Cloud Cover,
- Atmospheric Pressure,
- And Temperature.

What stands out in recent forecasting work? Renewable generation predictions draw the most AI study now (Benti et al., 2023; Krishnamurthy, 2024). Why does this area attract such focus? Power from renewables reacts sharply to weather shifts - tiny differences bring big swings in output. Because of these complex patterns, older straight-line methods often fail. Nonlinear dynamics make accurate modeling tough without newer tools.

Looking ahead, generation forecasts split based on time span. Different ranges bring distinct methods into play. Short-term views rely heavily on recent patterns. Medium spans adjust for shifting conditions more deliberately. Long-range estimates lean on broader trends instead. Each horizon shapes how predictions are built

- **Very Short Term Forecasting:** This type of prediction supports immediate balance within power networks. Though brief in scope, it plays a key role maintaining stability across electrical systems.
- **Short Term Forecasting:** Ranging from hours up to a day ahead - informs decisions in power dispatch alongside energy trading. While immediate needs shape grid operations, market activity relies on these estimates to balance supply with demand trends.

- **Medium-term forecasting:** Weeks to months, used for maintenance and resource planning.
- **Long Term Forecasting:** Predictions guide choices on roads and rules. Infrastructure plans rely on estimates made long before construction begins. Policy directions emerge from outlooks shaped by distant timelines.

Studies lately show machine learning works well for predicting near-future renewable output (Benti et al., 2023) - its strength lies in capturing quick shifts over time (Wang et al., 2023).

Load Forecasting

Forecasting electricity needs ahead of time forms the core idea behind load prediction. Long before solar and wind power spread worldwide, this type of projection ruled the field. Shifting habits in energy use come from a mix of factors that shape how much is drawn at any moment

- Weather conditions,
- Industrial activity,
- Population density,
- Economic growth,
- And behavioral trends.

Unlike predictions for renewable energy output, forecasts of electricity demand tend to be steadier, since people's usage habits repeat over time. Because of this regularity, older techniques like ARIMA models and smoothed trends delivered adequate results across numerous power networks in past decades.

Even so, rising use of electricity, growth in electric cars, spread of local power systems, and smarter grids have made demand shifts more unpredictable. As a result, algorithms based on machine learning now play a larger role in predicting loads (Chen et al., 2025; Jain et al., 2024) by uncovering subtle trends and evolving usage habits.

Electricity Price Forecasting

What makes predicting electricity prices so tricky? The answer lies in how power markets react - not just to supply and usage patterns, yet heavily influenced by shifts in fuel costs. Grid bottlenecks play a role too, alongside the tactics companies use when placing bids. Surprises on the global political stage often ripple through pricing models. Even policy changes introduced by regulators can reshape forecasts overnight.

Unlike physical demand or generation, electricity prices can exhibit:

- Spikes,
- Sudden volatility,
- Non-stationarity (changing statistical properties over time),
- Extreme reactions often follow sudden changes without warning.

Given these complications, researchers often find one-size-fits-all models fail to outperform others uniformly in every power market. As a result, blended artificial intelligence methods- linking traditional data breakdowns with algorithmic learning or neural networks - are gaining wider traction (Yandar et al., 2024; Shahverdi et al., 2025).

Traditional Statistical Forecasting Methods

Back then, predictions about energy needs leaned heavily on number patterns tracked across months or years. What you're looking at is data points lined up one after another, like snapshots taken regularly. These techniques try to spot repeating behaviors - maybe a steady climb, yearly cycles, or echoes from earlier numbers shaping today's figures.

Arima And Other Similar Models

What many turn to first is ARIMA, short for AutoRegressive Integrated Moving Average. This approach weaves together elements like

- Past values, differencing, along with error corrections
- Autoregression (using past observations),
- Differencing (removing trends),

Yet smoothing techniques track historical blips. Still, they rely on prior miscalculations. Even so, trends emerge from old gaps.

Patterns that repeat with the seasons get woven into SARIMA, building on earlier methods. These models became popular because:

- Math opens clear paths to understanding them
- Computationally inexpensive,
- Yet fairly useful when conditions stay steady and past trends repeat themselves.

Weather patterns twist in ways numbers often fail to capture, especially when predicting renewables. Lately, a string of recent papers found ARIMA methods lag behind during shifting skies (Benti et al., 2023; Krishnamurthy, 2024) - wind and sun forecasts go off track fast.

Advantages of Traditional Methods

Traditional methods still offer several important strengths:

- Low computational requirements.
- Faster training times.
- Easier interpretability.
- Lower implementation costs.
- Strong performance in stable and low-variability systems.

Even now, such benefits matter most where internet access is spotty or power grids are small. What counts shows up when systems lack speed, size, or support. Where tech runs thin, these strengths still hold ground.

Limitations of Traditional Methods

Even so, traditional methods come up short in more ways than one:

- hard to pin down patterns that twist and turn, yet often slip through standard methods meant for straight lines
- limited adaptability,
- weaker performance under volatile conditions,
- low precision shows up when predicting fast-changing renewable energy patterns.

When more renewables are used, such flaws start to matter far more.

The Rise of Machine Learning Forecasting

From raw data, machine learning pulls out repeating shapes in behavior - no step-by-step coding needed. Where statistical

methods rely on rigid rules set ahead of time, these tools shift and adapt on their own.

Lately, computers that learn from data have edged out old methods in predicting energy trends. While statistics once ruled, smarter models now handle the forecasts.

Common Machine Learning Models

The most widely used ML forecasting models include:

Model	Definition	Key Strengths
Random Forest	Ensemble of decision trees	Handles nonlinear data well
Support Vector Machine (SVM)	Separates patterns using optimal boundaries	Effective with smaller datasets
XGBoost / Gradient Boosting	Sequential error-correcting tree model	High predictive accuracy
K-Nearest Neighbors (KNN)	Predicts using similar historical cases	Simplicity

What helped these models catch on was their ability to beat older statistical techniques across most renewable prediction tasks. Yet unlike deep learning systems, they demand far fewer computing resources to run effectively.

Why Machine Learning Improved Forecasting

Machine learning models improved forecasting mainly because they can:

- Integrate large numbers of variables,
- Model nonlinear interactions,
- When things shift, they adjust more easily
- Working with data that has many dimensions happens here instead.

Take wind energy. It relies not just on how fast the wind blows, but also roughness in airflow, thickness of air, location features. Where standard math tools fall short spotting complex links among them, algorithms trained on data detect subtle patterns without being told where to look.

Most recent comparisons often show machine learning methods beating ARIMA models (Jain et al., 2024; Chen et al., 2025) on error scores like MAE and RMSE - especially when predicting wind patterns. Though older techniques still linger, newer systems tend to cut down mistakes more reliably across tests run lately.

Limitations in Machine Learning Systems

Even so, traditional machine learning systems still face limitations:

- They often require manual feature engineering,
- Struggle with long temporal dependencies,
- Yet they might miss how things follow one after another.

Creating features means adjusting inputs by hand so models work better. It takes a lot of time, often relying heavily on experience. Sometimes small changes make a difference - other times, nothing seems to help.

Flaws like these opened the door wide when deep learning started catching on around 2018 to 2020.

Deep learning moves toward neural architectures

What sets deep learning apart is its role as a pivotal shift in how researchers approach energy prediction today. Built on stacked layers, these neural systems detect complex structures within vast amounts of data without manual intervention. Though subtle at first glance, their ability to adapt through exposure reshapes accuracy over time.

LSTM and recurrent neural networks

LSTM networks, within deep learning approaches, took a leading role in studies on energy prediction. Though several techniques exist, these models stood out across much of the research.

LSTM belongs to the family of Recurrent Neural Networks, handling sequences by carrying forward information from earlier steps. Because of this trait, predictions over time-based data become more accurate. What sets it apart is how past inputs influence current outputs through internal gating mechanisms. Performance improves notably when patterns unfold gradually across moments.

What makes LSTM stand out lies in how it handles extended sequences over time

- Weather conditions from previous hours,
- Recurring consumption patterns,
- Delayed renewable responses,
- And sequential fluctuations.

When weather conditions shift unpredictably, LSTM models tend to forecast better than older methods(Benti et al., 2023). Though renewables rely on changeable climates, these networks adapt more effectively over brief periods.

CNN-LSTM Hybrid Models

Researchers increasingly developed hybrid architectures combining:

- CNNs (Convolutional Neural Networks), which identify spatial or local patterns,
- Using LSTMs helps identify patterns over time.

These hybrid systems became especially useful in:

- Solar forecasting using satellite/cloud imagery,
- Wind forecasting across multiple turbines,
- Multi-source prediction setups exist alongside varied data inputs.

Transformer Models

Lately, research into forecasting has begun exploring Transformer designs. These were first built for handling human language but rely heavily on attention processes - methods enabling them to highlight key sections within a sequence of data inputs.

Transformers are advantageous because they:

- Working on multiple pieces at once, computation moves faster through shared effort across units instead of step-by-step handling
- Processing extended data streams with improved effectiveness
- They perform efficiently even when handling substantial amounts of data.

Still, these methods demand heavy computation along with large datasets.

Statistical Machine Learning Deep Learning Compared

Starting with statistics, forecasts gradually shifted toward machine learning, eventually embracing deep learning - a progression shaped by balancing ease of use against precision and computing demands. While simpler models remain interpretable, their predictive power often falls short compared to complex architectures requiring heavy resources. Each step forward trades straightforward logic for better results at the cost of increased processing needs. Simplicity fades as systems grow more intricate, yet gains in accuracy justify the burden for many applications. This movement traces an ongoing compromise among clarity, performance, and technical requirements.

Traditional Statistical Models

- Cheapest,
- Most interpretable,
- Computationally efficient,
- Yet most fragile when facing unpredictable shifts in renewable sources.

Classical Machine Learning

- Improved nonlinear modeling,
- Moderate computational cost,
- Despite its limitations, the system shows reliable results when handling organized information
- Yet relies heavily on how features are crafted.

Deep Learning

- Among numerous analyses, predictions here tend to align most closely with outcomes
- Most capable when capturing patterns that unfold step by step,
- With little need for handcrafted features,
- Yet demands the most intense processing power alongside substantial hardware resources.

Most times, better predictions come at a price few stop to calculate. Though a neural network might reduce forecast errors by two or three points, its added complexity can inflate expenses in ways numbers alone won't show. So whether it's truly superior rests less on accuracy metrics, more on how systems run and budgets stretch in real settings.

Table 1: Comparative Overview of Forecasting Approaches

Approach	Main Strength	Main Weakness	Computational Cost	Typical Use
ARIMA/SARIMA	Simple, interpretable	Weak nonlinear modeling	Low	Stable load forecasting
Random Forest	Handles nonlinear data	Limited sequential memory	Medium	Wind/load forecasting
SVM	Effective with small datasets	Scaling limitations	Medium	Small-medium forecasting tasks
XGBoost	High predictive accuracy	Requires tuning	Medium-High	Structured forecasting problems
LSTM	Captures temporal dependencies	Expensive training	High	Renewable forecasting
CNN-LSTM	Captures spatial + temporal patterns	Complex architecture	High	Solar/wind forecasting
Transformer	Long-sequence modeling	Very compute-intensive	Very High	Advanced large-scale forecasting

Key Insights from Early Literature

Looking across the studies, certain trends appear repeatedly:

- Renewable integration is the main driver behind AI forecasting growth.
- AI approaches generally outperform statistical baselines in volatile renewable environments.
- Recently, deep learning has taken center stage in academic writing.
- Practical application often takes priority in grey literature, where usability outweighs strict precision.
- Whether better predictions cut actual operating expenses by more than it takes to set them up decides economic advantage.

Building on these ideas, the following part looks closely at a key factor behind accurate forecasts - how good the data is along with how features are built. What matters here isn't just volume but how well information is shaped. Often overlooked, preparation affects outcomes more than models do. Details like consistency, relevance, and structure play quiet yet powerful roles. Behind every strong prediction lies careful attention to inputs. How values are selected, cleaned, transformed shapes what comes out. Errors introduced early tend to spread silently forward. Patterns hidden in raw numbers emerge only when handled right. Success shows up less in algorithms, more in groundwork.

Data Quality and Features Affect Forecast Accuracy

Surprisingly often, studies show forecast precision relies less on algorithms alone (Chen et al., 2025; Jain et al., 2024) - more on how clean the input information happens to be. Even though machine learning methods keep evolving into complex systems, experts point out flawed datasets weaken predictions regardless of design strength. Because of this pattern, trustworthy inputs now shape outcomes just as much as code does in power output, demand, and market cost estimates.

Although energy forecasts depend on accurate, complete, consistent, and up-to-date data, flawed inputs often go unnoticed. Wind speed, sunlight levels - defined as solar energy hitting a surface - cloud density, moisture in the air, and heat measurements shape predictions for renewable output. When estimating electricity use, factors like dates, market trends, how people consume power, and past usage emerge as relevant. Since artificial intelligence adapts based on what it sees, poor-quality numbers tend to distort outcomes behind the scenes. Predictive strength weakens when foundational records carry mistakes.

Feature engineering draws attention across many scholarly works (Chen et al., 2025) - researchers describe it as shaping or choosing variables so models detect relevant structures more effectively. Though methods like Random Forests rely heavily on thoughtfully built inputs, such as past consumption figures, rolling summaries, time-based markers, or climate shifts, newer approaches differ. Architectures rooted in deep learning (Benti et al., 2023), including LSTMs and Transformers, pull out insights from unprocessed sequences without human-guided adjustments. Even then, their success still ties closely to how clean and informative the incoming data appears. Evidence repeatedly points to one idea: better

predictions come just as much from sharper weather estimates and refined datasets as they do from advanced algorithms.

When looking at scholarly versus grey literature, a clear distinction appears. While academic research tends to work with neat, well-organized data, its main aim is refining computational methods. On the other hand, industry publications highlight real-world hurdles in gathering and aligning information from diverse platforms. Problems like gaps in sensor output, mismatched database formats, slow transfer speeds, and weak oversight policies are common among utility providers. Because of these conditions, non-academic sources typically point to weaknesses in data systems - rather than model choice - as the key obstacle preventing effective AI implementation.

This difference matters most in places like India, where growth in green power happens at high speed. While spending on intelligent networks, modern meters, internet-connected detectors, and centers managing renewables boosts how much information exists, the reliability of that data along with smooth tech links varies widely by location. Under these conditions, gains from artificial intelligence predictions rely just as heavily on foundational data systems as they do on the design of predictive tools.

Forecasting ability tends to depend heavily on how good the underlying information happens to be. Though artificial intelligence may help uncover hidden trends within datasets, flawed inputs still limit what any model can achieve. Reliable evidence forms the foundation - without it, even sophisticated methods struggle. Experts across academia and practice agree on this point: trustworthy results start with solid data. Quality does not just support prediction systems - it enables them.

Classical Machine Learning Compared to Deep Learning When Predicting Outcomes?

Energy forecasting's evolution brings a persistent debate: do more intricate AI systems truly improve predictions? During the last ten years, scholars moved step by step away from classic statistics, embracing first machine learning, then later deep learning frameworks. Superior accuracy frequently appears with deep networks - yet such performance comes paired with heavier computing demands and higher expenses. These trade-offs spark scrutiny on whether added layers of sophistication deliver proportional benefits. Thus emerges a central theme in current studies - the balance between conventional ML methods and deeper neural approaches.

What made Random Forest, SVM, and XGBoost stand out (Yandar et al., 2024; Abisoye, 2024) was their ability to move beyond older statistical techniques without demanding excessive computing resources. Because of how they handle curved or twisted data shapes, these tools fit naturally into predictions for renewables, where output shifts under tangled influences like weather swings, time-of-year effects, and equipment behavior. When tested across many cases, algorithms from the machine learning family tend to beat classic options - ARIMA being one example - especially when dealing with sun- or wind-based systems, since straight-line math rarely captures real-world variability accurately.

Though traditional machine learning struggles with certain constraints, deep learning structures arose to address them. Because they process sequences effectively, LSTM networks handle long-term patterns across time intervals. GRUs function similarly but rely on simplified control mechanisms to track evolving inputs. Meanwhile, models built on Transformers detect distant connections in data without stepping through each point sequentially. These methods excel where past conditions shape present outcomes. Since power production from wind and solar shifts gradually, earlier states influence later readings. Demand cycles also unfold progressively rather than changing at random. As a result, fixed-variable assumptions fall short when applied to dynamic energy systems.

Lately, research shows deep learning often wins (Benti et al., 2023; Krishnamurthy, 2024) in predicting short-term renewable energy output. Rather than older techniques, methods built on LSTM tend to show smaller mistakes when tested with measures like MAE, RMSE, and MAPE. In several cases, Transformers push results even further (Mamat, 2025; Moshtaghi et al., 2025) - spotting distant patterns in big data better than standard RNNs ever did. Because of these gains, studies now lean heavily toward deep learning for forecasts.

Yet superior predictions alone do not confirm broader advantages. One major drawback of deep learning lies in its heavy computational demands. Building extensive neural networks involves high-powered hardware, vast data volumes, extended training periods, alongside expert knowledge. On the other hand, traditional approaches like Random Forests or XGBoost deliver similar accuracy using far less computing effort. Such differences matter most when prediction tools run nonstop outside lab conditions.

Though rarely discussed in journals, real-world challenges emerge clearly in grey literature. Instead of chasing slight gains in prediction precision, industry analyses emphasize durability under load, upkeep demands, transparency, and running expenses. Even if method boosts forecast correctness slightly, its need for heavy computing resources might erase any monetary advantage. Because of this, power providers favor tools that blend decent accuracy with steady operation, clear logic, and affordable implementation.

Understanding how a model reaches its conclusions matters just as much as the results it produces. While traditional machine-learning methods tend to reveal their logic more clearly, deep neural networks usually do not. In tightly controlled energy sectors, being able to back up forecasts becomes essential - regulators, stakeholders, or planning teams may demand clear reasoning. Even if they deliver precise outcomes, deep learning systems frequently face skepticism due to opaque inner workings.

Deep learning tends to lead in prediction quality - especially when handling big data in tricky renewable energy forecasts. Yet classical methods frequently strike a better mix of precision, speed, and real-world usability. A noticeable compromise surfaces upon comparison. The research overall refuses to back any one-size-fits-all claim favoring deep

learning. Context shapes how well a model performs - not just its design.

Because of money-related factors, this difference matters a lot. When a deep-learning system offers little gain compared to standard machine learning yet demands far more computing power and effort, it might not justify the expense. On the flip side, in energy grids rich with renewables, where balancing loads gets expensive or forecast precision is tightly regulated, tiny gains in prediction quality can still save significant amounts. So whether to choose machine learning or deep learning turns out to hinge less on performance alone and more on cost-effectiveness. What works best isn't always the top-scoring algorithm, rather the one that brings highest usefulness for the resources spent.

Probabilistic Forecasting and Managing Uncertainty

Though AI-driven forecasts have grown more precise, experts now see a core flaw in classic methods: perfect certainty remains out of reach, no matter how sophisticated the system. Sudden shifts in weather pop up without warning. Machines break down at odd times. Shoppers alter their choices in unpredictable patterns. Because of this, newer studies moved past single-value predictions (Mayer et al., 2023), turning instead to methods that map the range of possible outcomes. Uncertainty gets measured now, not set aside.

A single expected number is what a point forecast delivers. Imagine predicting a solar plant will produce exactly 500 megawatts at midday tomorrow. Simple to grasp, yes - yet such estimates say nothing about uncertainty. Instead of one figure, probabilistic methods lay out multiple potential results. Each possibility comes paired with how likely it may occur. A forecast might not pick one number. Instead, it can suggest power output will probably stay within 470 MW and 530 MW, given a certain confidence. Such range estimates help leaders see possible outcomes across varying conditions. Understanding likelihoods becomes part of planning, not just central predictions.

Weather patterns shape how we predict power output from sun and wind sources. Because these sources rely on conditions beyond control, surprise variations happen often. Recent research suggests focusing only on shrinking mistakes misses part of the picture. Instead, showing likely ranges builds better decisions. A prediction modest in precision but clear about doubt might serve grids better. Exact numbers without context carry hidden dangers when storms shift plans.

Most studies back probabilistic forecasts (Mayer et al., 2023), especially when predicting power from renewables or market prices. Various strategies have emerged - some rely on prediction ranges, others on specific percentiles, Bayesian logic, or grouping several models together. Instead of one fixed number, combining outputs across different systems creates a spread of likely results. This mix tends to capture unpredictability better. It gives grid managers clearer insight when facing sudden shifts.

Still, industries have taken up probabilistic forecasting at a slower pace compared to how fast it spreads in scholarly work. While unofficial reports recognize the usefulness of knowing uncertainties, they also bring out real-world hurdles

in using such methods daily. Instead of clear-cut numbers, these forecasts give ranges - something energy providers and traders sometimes struggle to apply, especially since existing procedures expect one figure per prediction. Because infrastructure and rules evolved around fixed estimates, shifting mindset takes time. So progress outside labs tends to trail behind what researchers develop.

This gap sets apart academic work from informal sources in a notable way. While experts often see measuring uncertainty as key to better predictions, those in business care more about turning probability data into clear actions. A prediction adds little value unless users understand what to do when outcomes are uncertain.

One big plus of probabilistic forecasting? It helps make smarter money choices. When predictions miss the mark, expenses pile up - not just from error but from surprise swings in supply or demand. Because these forecasts show how likely various scenarios are, decision makers gain clarity on where to place backup resources. This shift often leads to lower costs when matching supply with actual use. As renewables take up more space in grids like India's, preparing for unknowns turns less optional. Over time, knowing probabilities could matter more than guessing exact values.

One key takeaway from current research is that progress in energy prediction might matter less than how well we handle unpredictability. Though single-value estimates still guide most daily operations, methods showing a range of possible outcomes are gaining ground - offering clearer insight into risks. With renewables spreading across power systems worldwide, knowing the limits of a forecast could soon weigh equally with predicting actual generation levels.

The Economic Impact of Using AI to Predict Energy Needs

Most research judges prediction tools by numbers like MAE, RMSE, or MAPE. Yet behind those figures sits a more pressing concern - does better forecasting actually pay for itself? That gap shapes much of the discussion around artificial intelligence in power demand predictions. Even if an AI system outperforms others on paper, its real-world impact might remain negligible. Because of this, linking forecast precision to actual financial outcomes becomes critical when weighing whether advanced methods are worth adopting.

Usually, studies on predictions work from the idea that fewer mistakes mean better results in practice. Although this holds true in many cases, the link between error size and performance isn't uniform across levels. Cutting forecast errors in half - from 10% down to 5% - can lead to meaningful cost savings; however, shrinking them just one point more might barely make a difference at all. Such patterns reflect what economists call diminishing returns: each extra effort brings less gain than the last. Hence, real-world impact hinges not solely on precision metrics but also on how forecasts guide decisions inside power networks.

Most obvious among financial gains from precise predictions? Lower balancing expenses (DeepMind, 2019; IEA, 2021). Keeping supply aligned with demand defines grid operation. Underestimating green output sometimes wastes

clean power - cutbacks follow, income drops. Should production exceed forecasts, authorities risk gaps needing costly real-time purchases. Overlooked surplus often means unused potential. Forecasting that hits closer to reality cuts down on waste, mainly because timing and delivery choices get better. Because of this, many research efforts point to lower balancing expenses as a key reason to fund forecast improvements.

A bonus comes through smaller needed reserves. When forecasts improve, guesswork drops. Operators keep backup power ready - just in case predictions miss sudden shifts. Yet wind and solar change output on their own rhythm. Heavier reliance on these sources usually means more spare capacity held waiting. With sharper prediction tools, confidence grows about what lies ahead. Less idle supply stays online as a result. Improving system efficiency often comes hand in hand with reduced operating expenses. Research papers regularly point to better reserve management (NREL, n.d.; IEA, 2021) as a key driver of financial gains - especially where wind energy plays a dominant role.

Forecasting shapes financial outcomes in electricity markets, too. Where competition exists, those who supply power rely on predictions when setting bids, planning output, or handling monetary uncertainty. Better forecasts often lead to sharper trades while lowering vulnerability to sudden price shifts. Price projections matter most here - mistakes tend to become real-dollar setbacks. Small upgrades in prediction accuracy, then, sometimes yield surprisingly large returns.

Still, weighing those advantages means also looking at what it takes to run AI tools. Because deep-learning systems demand heavy computing power, vast amounts of data, skilled personnel, along with continuous updates. Power companies might have to spend on cloud platforms, sensing devices, protection from cyber threats, plus staff education - just to get AI-based forecasts working well. Unofficial reports often point out these behind-the-scenes expenses, suggesting full lifetime costs matter more than just how precise a forecast appears. Seen this way, highest precision does not automatically mean best value.

What stands out is the difference between peer-reviewed research and industry publications. While university-based work measures success mainly by prediction accuracy, company-driven analysis asks if better forecasts lead to real-world gains. Instead of ranking algorithms, private sector reviews ask how changes affect costs, efficiency, or decision speed. Even a slight drop in error rate won't justify cost unless it shifts outcomes that matter to leadership. Because of this, practical guides stress factors like ease of rollout, long-term maintenance, and financial payoff over model complexity.

Renewable growth in India sharpens the link (IEA, 2021; CEA, n.d.; Prayas Energy Group, n.d.) between prediction quality and money at stake. Because forecasts shape how power is scheduled, balanced, and managed on grids, even small mistakes add up - penalties appear, systems slow down, backup needs grow. Where conditions shift fast, artificial intelligence may cut guesswork, steering choices with better clarity. Value emerges quietly, not from perfection, but from fewer surprises in daily operations. While some areas gain

greatly from improved predictions, others see little advantage due to differences in tech access, local needs, or how developed their markets are. Because of this, the financial upside of artificial intelligence in forecasts shifts based on setting - rarely applying equally everywhere.

What stands out in newer studies is the idea that forecast precision serves a purpose beyond itself. Instead of focusing purely on shrinking prediction errors, attention shifts toward real-world results - how forecasts influence business performance or financial efficiency. Success then depends less on statistical perfection and more on usefulness when guiding choices. Sometimes a simpler method, low in expense and effort to maintain, proves more beneficial even if its numbers are slightly off. High complexity often brings small gains in correctness while demanding much more in resources, making such approaches harder to justify.

Altogether, studies point toward meaningful financial gains when using AI-driven forecasts - especially where renewables dominate and grid balancing gets expensive. Still, such advantages aren't automatic. Whether they materialize hinges on how precisely models predict, what they cost to run, the setting they're used in, and how markets are structured. For this reason, research stops short of claiming AI methods outperform alternatives across every situation. Yet, economic benefit emerges only if better decisions bring gains larger than the expenses tied to building, running, and updating the AI forecast tool. One key takeaway from the analysis rests here - weighing costs against outcomes matters deeply when judging different forecasting methods.

The Deployment Gap Scholarly and Grey Literature Compared

What stands out clearly through this analysis is how differently academics and professionals judge forecasting tools. Even though improvement from artificial intelligence receives broad recognition across fields, measures of progress rarely align. In journals, precision and novel techniques take center stage. Meanwhile, non-academic reports emphasize usability, expansion potential, execution speed. That mismatch forms a noticeable space - between high-scoring models in papers and workable solutions onsite - one thrives in theory, the other must survive daily operations.

Most research aims to push forward how forecasts are made (Moshtaghi et al., 2025). Because of this, scholars spend time building more complex designs, testing different models against one another, while also chasing small gains in numbers like MAE, RMSE, or MAPE. Whether a fresh method beats older ones on standard data sets usually defines progress. Such priorities have sped up change - especially within deep learning, mixed prediction tools, and methods that estimate uncertainty. Though many newer papers show advances of just a couple percent, those results tend to be framed as meaningful steps ahead.

Forecasting, within grey literature, takes on a distinct angle. From the standpoint of consulting groups, utility providers, grid managers, and policy bodies, predictive models are rarely isolated tools but fit into larger operational frameworks. Instead of focusing solely on slight gains in precision, these actors prioritize consistent deployment, ease of upkeep, and alignment with current workflows. Attention shifts - security

concerns, system compatibility, staff proficiency, adherence to regulations, and preparedness of physical systems weigh just as heavily as how correct predictions turn out. What matters grows beyond numbers when real-world constraints shape expectations.

What stands out is how views split over deep learning uptake (Shahverdi et al., 2025). While studies highlight complex neural networks as leading tools for predictions - thanks to strong accuracy records - their real-world application tells another story. Practical guides point to hurdles like heavy computing needs, ongoing model updates, niche skills, and vast volumes of clean data. These demands make widespread deployment less straightforward than theory suggests. Most groups find these demands tough to meet, slowing wide adoption. Still, power companies might stick with basic machine learning methods when those work well enough while being cheaper and easier to manage.

What sets these groups apart also ties into clarity. While researchers tend to value precision - despite opaque mechanisms - those applying models in real-world settings usually demand insight into decision logic. Because forecasts guide budgeting, legal adherence, and workflow design, a system's inner workings matter just as much as its output. Even if it performs marginally better, an impenetrable method might face resistance when users struggle to trace prediction pathways. That unease fuels momentum behind interpretable machine learning methods designed to maintain accuracy while revealing process details.

Even though they differ, scholarly work and informal publications rarely oppose each other. Often, their focus falls on separate phases of how technology evolves. While universities probe the edges of technical potential, companies test whether ideas function reliably in real conditions. New solutions typically emerge in labs, then slowly shift toward field use. Moving from concept to implementation tends to stretch over long periods, especially within essential systems like power networks, where safety margins leave little room for error.

Most times, economic returns from AI forecasting depend less on model precision and more on real-world application. Although research papers highlight gains in predictive accuracy, field reports stress execution as the deciding factor. Even highly refined systems might underperform once moved outside controlled environments. Integration challenges or upkeep demands often limit their usefulness in day-to-day operations. Simpler alternatives sometimes outshine complex ones, simply because they fit smoothly into existing processes. What works in theory does not always hold up when put into practice.

One key takeaway stands out when examining this topic - forecasting skill by itself cannot measure how useful an AI method really is. Better algorithms do not necessarily lead to better outcomes in practice. What works well often depends on a mix of precision, clarity, ease of scaling, and expense. Looking across peer-reviewed studies and industry reports reveals a pattern: progress will hinge less on breakthrough models than on connecting lab ideas to actual use. Success comes from aligning technical gains with practical demands.

3. Synthesis of Themes and Key Findings

Looking beyond individual cases, one trend stands clear. Where renewable output shifts with complex weather patterns, artificial intelligence tends to beat older statistical techniques. Instead of relying solely on model complexity, results hinge largely on how clean and rich the input data really is. Rather than always chasing top accuracy, many find simpler algorithms deliver better value when weighing expense against clarity. Though deep networks sometimes lead in precision, they do not automatically justify their added demands.

Although forecasts once focused on precise predictions, recent work emphasizes handling uncertainty more directly. Rather than aiming for exact outcomes, researchers now prioritize measuring unpredictability through probability-based methods. Such approaches gain importance in energy grids packed with renewables, given their inherent fluctuations. Facing these shifts, planners treat variation not as noise but as a core factor in system design.

What matters most is how the data points to a limited financial upside from AI predictions. Value appears only if better forecasts lead to stronger choices in practice, enough to cover expenses tied to rollout. So, shifts arise not from proving AI beats older methods - it usually does - but from seeing if such improvements matter where work actually happens.

One key takeaway stands out when comparing peer-reviewed studies with informal reports: progress in theory often misses real-world use. Though researchers chase higher accuracy in predictions, businesses prioritize solutions that deliver clear benefits. This divide may shape much of what comes next for tools designed to forecast energy demand.

4. Future Research Directions

Despite progress in AI-driven energy forecasts during recent years, key questions still go unanswered. Success often hinges too heavily on prediction precision alone - a habit that overlooks broader implications. While many papers rely on standard error measures like MAE, RMSE, or MAPE, they rarely link better scores to real-world financial gains. Instead of fixating on numbers, upcoming work might benefit from weighing how well models reduce expenses tied to grid balance, backup capacity, trading outcomes, and capital efficiency. Shifting focus this way could reveal what truly works in practice.

Understanding how artificial intelligence reaches its conclusions remains a challenge that needs more study. Though models grow ever smarter, seeing inside their decision-making process gets harder. When people cannot follow the logic behind an output, doubt creeps in - among those who run systems, oversee rules, or provide funding. Rather than chasing only higher precision, newer methods might place equal value on clarity. Seeing which inputs shape outcomes helps users grasp results, even when complexity stays high. Accuracy need not vanish just because insight improves.

One reason progress continues slowly lies in how few teams apply uncertainty data during live operations. Although evidence supports better outcomes with probabilistic methods, adoption lags behind theory. Instead of treating forecasts as passive inputs, future efforts could tie predictions more closely to responses. Questions remain about fitting these models into trading tactics, control workflows, or infrastructure planning. What happens next depends on bridging the gap between statistical output and actual choices.

One path worth exploring focuses on building forecasting tools that demand less computing power. Even though complex systems like Transformers perform well, heavy resource needs could slow real-world use - especially where infrastructure is limited. Future research should therefore focus on identifying forecasting approaches that achieve an effective balance between forecasting accuracy, implementation cost, and deployment scale. While simpler models can reduce hardware and electricity requirements, larger-scale deployments may lower the cost per forecast by spreading computational resources across wider electricity networks. Determining the optimal balance between scale, cost, and accuracy will therefore be an important direction for future research.

Looking ahead, emerging energy systems in developing economies deserve greater scrutiny. While most studies focus on North America, Europe, or China, countries such as India are experiencing some of the world's fastest growth in renewable energy capacity. Future research should therefore examine how forecasting models perform within rapidly expanding grids, uneven digital infrastructure, and evolving regulatory environments. In addition, researchers may explore the role of AI forecasting in smaller-scale renewable systems, including rooftop solar installations, community energy projects, and microgrids. Such developments could contribute to the wider democratization of forecasting technologies, allowing households, businesses, and local communities to benefit from forecasting tools that have traditionally been available only to large utilities. These trends may ultimately reshape both the economics and accessibility of renewable energy management.

5. Limitations

Looking ahead, emerging energy systems in developing economies deserve greater scrutiny. While most studies focus on North America, Europe, or China, countries such as India are experiencing some of the world's fastest growth in renewable energy capacity. Future research should therefore examine how forecasting models perform within rapidly expanding grids, uneven digital infrastructure, and evolving regulatory environments. In addition, researchers may explore the role of AI forecasting in smaller-scale renewable systems, including rooftop solar installations, community energy projects, and microgrids. Such developments could contribute to the wider democratization of forecasting technologies, allowing households, businesses, and local communities to benefit from forecasting tools that have traditionally been available only to large utilities. These trends may ultimately reshape both the economics and accessibility of renewable energy management (Shahverdi et al., 2025; CEEW, n.d.; NSGM, 2017).

Now shaping up fast, artificial intelligence keeps shifting beneath our feet. With fresh forecast designs appearing alongside novel data sets and testing methods, earlier judgments might need updating once more results arrive. Especially true for rising techniques like Transformer-driven prediction systems or broad-purpose foundation models - how they perform over time and behave in real settings still unfolds through ongoing study.

Looking at different studies is tricky, since each uses distinct forecasting methods, data sources, locations, or ways of measuring success. One setting might show strong performance - another reveals weaker outcomes under similar claims. Because of these shifts in context, drawing firm parallels becomes problematic without deeper scrutiny.

Though academic sources typically outline methods clearly, industry documents sometimes lack such transparency - making verification harder when assessing efficiency or financial impact. Still, blending peer-reviewed studies with real-world data paints a fuller picture of what AI can - and cannot - deliver in energy prediction tasks. Verification gaps exist, true, yet diverse inputs help capture practical hurdles alongside theoretical potential.

6. Conclusion

Looking into today's AI-based tools for forecasting energy needs was the main goal. Did these approaches do better than older versions? That mattered most. Findings emerged by studying scholarly articles together with less formal documentation. Where one method gained an edge, precision often tipped the balance. Tricky setups tended to reveal deeper challenges during real-world use. Over time, simpler models tended to demand fewer updates. Even with possible advantages, financial factors shaped how widely they were adopted. During tests in actual conditions, odd trends began appearing. Their economic value relied strongly on specific settings and size of operation. Results suggested better performance in certain cases, not everywhere.

Unexpected improvements appear once artificial intelligence takes on energy forecasting. When traditional statistics falter - particularly around erratic behavior - methods powered by AI deliver more accurate predictions for supply, usage, and costs (Benti et al., 2023; Abisoye, 2024). Renewable sources bring chaos through unstable weather signals; oddly enough, such disorder plays directly into machine learning's strengths. Instead of fixed equations, modern systems like neural nets detect subtle patterns buried deep inside continuous data flows. As conditions on the electrical network shift unpredictably, their ability to adjust quietly outperforms legacy solutions. Few can overlook that advantage today.

Even so, findings suggest that accurate predictions depend on more than just advanced models. What counts is clean input - also, picking meaningful indicators, paired with reliable tech setups. Papers and field notes alike confirm a pattern: top-tier algorithms stumble if information lacks quality or consistency (Chen et al., 2025; Jain et al., 2024). When details are messy or unevenly managed, performance drops regardless of method sophistication. For this reason, progress in tracking

systems, updated power grids, and tighter review processes weighs as much as tweaking forecast designs.

Deep learning may score well on forecasts, yet requires massive datasets, powerful hardware, often specialized knowledge (Benti et al., 2023; Shahverdi et al., 2025). Impressive as they seem, such models face competition from basic techniques - these frequently bring transparency, faster deployment, reduced costs (Yandar et al., 2024; Krishnamurthy, 2024). The setting guides the decision; raw accuracy fails to capture suitability. Factors like operational demands, accessible resources, and staff background shape each selection. Higher metrics? Not necessarily smarter results.

Forecasting based on probabilities is becoming more common, showing that people are starting to accept uncertainty as part of energy systems. Moving beyond single-point predictions, new methods work by capturing ranges of possible outcomes - this supports better decisions when future conditions are unclear. With renewable sources playing a larger role worldwide, ignoring these fluctuations might soon seem unwise. What began as an uncommon approach has shifted into a central feature, guiding how projections adapt throughout electricity grids.

Noticeable benefits emerge only under certain conditions when using AI-based forecasting. While these systems often outperform traditional approaches in accuracy, that edge does not automatically translate into profit improvements. Success depends on factors such as operational costs, reserve capacity demands, regulatory frameworks, renewable supply levels, and initial investment size. Where inaccurate estimates carry serious consequences - be they logistical setbacks or financial penalties - even slight advances in prediction quality may result in significant advantages. In contrast, similar enhancements might barely register where risks are low or margins wide. Even when pointing toward comparable outcomes, these source kinds diverge in what they value most. While academic work highlights accuracy alongside methodological advances, corporate analyses fixate on steady outputs, growth prospects, then margins - repeatedly. Bridging such contrasting views remains a persistent challenge in improving energy demand forecasts.

Most of the time, focusing only on lower mistakes in AI-based energy forecasts overlooks what actually counts. It is a clearer impact on choices that defines value - how consistently those predictions guide actions with measurable cost benefits. Success tends to appear in places rich with renewable sources, given solid data exists and there's clear incentive to respond quickly. As global power systems evolve toward sustainability and automation, dependence on forecasting tools could increase steadily. Their long-term importance will depend less on exact numbers, more on lasting influence after the forecast fades.

References

- [1] Abisoye, B. O. (2024). A survey of artificial intelligence methods for renewable energy forecasting. *Renewable Energy Focus*.

- [2] Benti, N. E., Chaka, M. D., & Semie, A. G. (2023). Forecasting renewable energy generation with machine learning and deep learning: Current advances and future prospects. *Sustainability*, 15(9), 7087. <https://doi.org/10.3390/su15097087>
- [3] Chen, Y., et al. (2025). A review of machine learning techniques for building electrical energy consumption prediction. *Energy and AI*. <https://www.sciencedirect.com/science/article/pii/S2666546825000503>
- [4] Jain, A., et al. (2024). Evaluation of electrical load demand forecasting using various machine learning algorithms. *Frontiers in Energy Research*. <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1408119/full>
- [5] Krishnamurthy, S. (2024). Artificial intelligence-based forecasting models for generation and load: A review. *Energy for Sustainable Development*.
- [6] Mamat, R. (2025). Emerging role of generative AI in renewable energy forecasting and optimization. *Energy Reports*.
- [7] Mayer, M. J., et al. (2023). Probabilistic modeling of future electricity systems with high renewable shares using artificial neural networks. *Applied Energy*. <https://www.sciencedirect.com/science/article/pii/S0306261923001654>
- [8] Moshtaghi, P., et al. (2025). A comprehensive review of artificial intelligence applications in power systems (2015–2024). <https://www.sciencedirect.com/science/article/pii/S2666188825002060>
- [9] Shahverdi, N., et al. (2025). A systematic review of machine learning and AI for sustainable energy systems. *Renewable and Sustainable Energy Reviews*.
- [10] Yandar, K. L., et al. (2024). Machine learning techniques in electric power prediction: A systematic literature review. *Journal of Electrical Systems*.
- [11] Grey Literature and Industry Sources
- [12] Council on Energy, Environment and Water (CEEW). (n.d.). Clean energy reports. <https://www.ceew.in>
- [13] Central Electricity Authority (CEA). (n.d.). Load generation balance reports. <https://cea.nic.in>
- [14] DeepMind. (2019). Machine learning can boost the value of wind energy. <https://deepmind.google/blog/machine-learning-can-boost-the-value-of-wind-energy/>
- [15] International Energy Agency (IEA). (2021). Renewables integration in India. <https://iea.blob.core.windows.net/assets/7b6bf9e6-4d69-466c-8069-bdd26b3e9ed1/RenewablesIntegrationinIndia2021.pdf>
- [16] National Renewable Energy Laboratory (NREL). (n.d.). India renewable integration study. <https://www.nrel.gov/analysis/india-renewable-integration-study>
- [17] National Smart Grid Mission (NSGM). (n.d.). Smart grid overview. <https://www.nsgm.gov.in/en/smart-grid>
- [18] National Smart Grid Mission (NSGM). (2022). Mission factsheet. <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2022/jun/doc202262166701.pdf>
- [19] National Smart Grid Mission (NSGM). (2017). India smart grid country report (Mission Innovation). <https://www.nsgm.gov.in/sites/default/files/India-Country-Report-Mission-Innovation-C1-Smart-Grids-June-2017.pdf>
- [20] Prayas Energy Group. (n.d.). Grid integration of renewables in India. <https://energy.prayaspune.org/our-work/research-report/grid-integration-of-renewables-in-india>
- [21] McKinsey & Company. (2024). The state of AI in early 2024: Gen AI adoption spikes and starts to generate value. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- [22] Deloitte. (2023). Tech trends 2024. Deloitte Insights. <https://www2.deloitte.com/us/en/insights/focus/tech-trends.html>