PHYSICS-INFORMED GENERATIVE ARTIFICIAL INTELLIGENCE FOR MODELING, PREDICTION, AND OPTIMIZATION OF SMART ELECTRICAL SYSTEMS

[1]Saud Zaman Khan, [2]Mirza Alizar Sami, [3]Sonia Farah, [4]Sultan Abbas, [5]Yasir Ali
[1]Computer Science Bachelor, Iqra University, Islamabad
[2]Department of Data Science, National University of Science and Technology MISIS (Moscow - Russia).
[3]Department of Telecommunication Engineering, University of Engineering and Technology, Taxila.
[4]Computational & Theoretical Physics Research Centre, Khawja Fareed University of Engineering and Information Technology, Rahim Yar Khan, Pakistan.
[5]Department of Electronics Engineering, NED University of Engineering and Technology Karachi.
[1]Saudzaman96@gmail.com, [2]alii.zar090@gmail.com, [3]engr.soniafarah@gmail.com, [4]sultanabbas8285@gmail.com, [5]yasirali014412@gmail.com

DOI: https://doi.org/10.5281/zenodo.22201125

Article History
Received on: 28 Feb, 2026
Accepted on: 26 March, 2026
Published on 29 March, 2026

Corresponding Author

Abstract
Purely data-driven artificial intelligence models for smart electrical systems suffer from the need for large amounts of labeled data, the lack of guaranties on physically meaningful outputs, and the tendency to produce predictions and scenarios that violate the laws governing power networks . As grids are incorporating larger shares of distributed and uncertain resources, these shortcomings are becoming increasingly consequential. In this work, we propose a physics-informed generative artificial intelligence framework that unifies the modeling, prediction and optimization of smart electrical systems in a single architecture. The framework integrates a conditional denoising diffusion model that captures the joint uncertainty of loads, renewables, and network states, with a differentiable physics layer that embeds the underlying electrical laws, i.e., AC power-flow equations, voltage and branch limits, and system power balance, directly into the learning objective. The training is based on a composite loss function that combines a data loss, a physics constraint violation term and a regularization term with tunable weighting factors. This allows the model to generate scenarios that are both statistically faithful and physically consistent. The generated scenario ensembles serve as a probabilistic prediction engine and a scenario based stochastic optimization module. The physics layer verifies and projects the outputs on the feasible manifold during inference. Key findings show that the physics-informed formulation offers a beneficial inductive bias that improves generalization in data-scarce regimes, substantially reduces constraint violations, and limits physically implausible drift in long-horizon forecasts. The generative module produces full joint scenario ensembles and effective data augmentation, and the optimization module reaches near-optimal decisions at a significantly lower computational cost than traditional iterative solvers. Ablation studies show that these physics constraints and the generative mechanism are complementary to each other in providing gains. The framework is of practical value for day-ahead and

intraday probabilistic forecasting, renewable scenario generation, synthetic data production under privacy constraints, and dispatch and market operations. Future work will include foundation model architectures, explicit feasibility guaranties, multi-energy extensions and real-time digital twin integration.
Keywords: Physics-informed AI; generative AI; smart electrical systems; diffusion models; probabilistic forecasting; scenario generation; power system optimization; constraint satisfaction; deep learning

1. Introduction

Modern electrical power systems are undergoing a profound transformation driven by global decarbonization commitments, the large-scale integration of renewable energy sources, and the widespread deployment of distributed energy resources, electric vehicles, and advanced sensing and communication infrastructure. This evolution toward the smart grid has dramatically increased the operational complexity of electrical systems: generation has become stochastic and spatially dispersed, load patterns are increasingly dynamic and uncertain, and control decisions must now be made across faster timescales and heterogeneous network layers. At the same time, the proliferation of smart meters, phasor measurement units, and Internet-of-Things devices has produced an unprecedented volume of operational data. Consequently, classical model-based methods for modeling, prediction, and optimization—while valuable—are often stretched to their limits by the scale, nonlinearity, and uncertainty of modern smart electrical systems, motivating new computational paradigms that can learn from data while remaining grounded in physical reality.

Over the past two decades, advances in artificial intelligence and deep learning have offered powerful new tools for addressing these challenges. Neural networks have achieved remarkable performance in load and renewable-generation forecasting, fault diagnosis, stability assessment, and energy management. More recently, generative AI—encompassing generative adversarial networks, variational autoencoders, and diffusion models—has extended these capabilities by learning the underlying probability distributions of system data. These models can synthesize realistic operating scenarios, augment scarce datasets, capture the stochastic behavior of renewable generation and demand, and support simulation-driven planning and optimization. However, purely data-driven approaches suffer from well-documented limitations. They typically require large volumes of high-quality training data, which are often unavailable for emerging grid configurations; their outputs can be physically inconsistent, violating conservation laws, power-flow equations, or operational limits; and they generalize poorly to conditions outside the training distribution. The black-box nature of these models further hinders trust and explainability in safety-critical grid operation.

Figure 1: Overall concept of Physics-Informed Generative AI for Smart Electrical Systems.

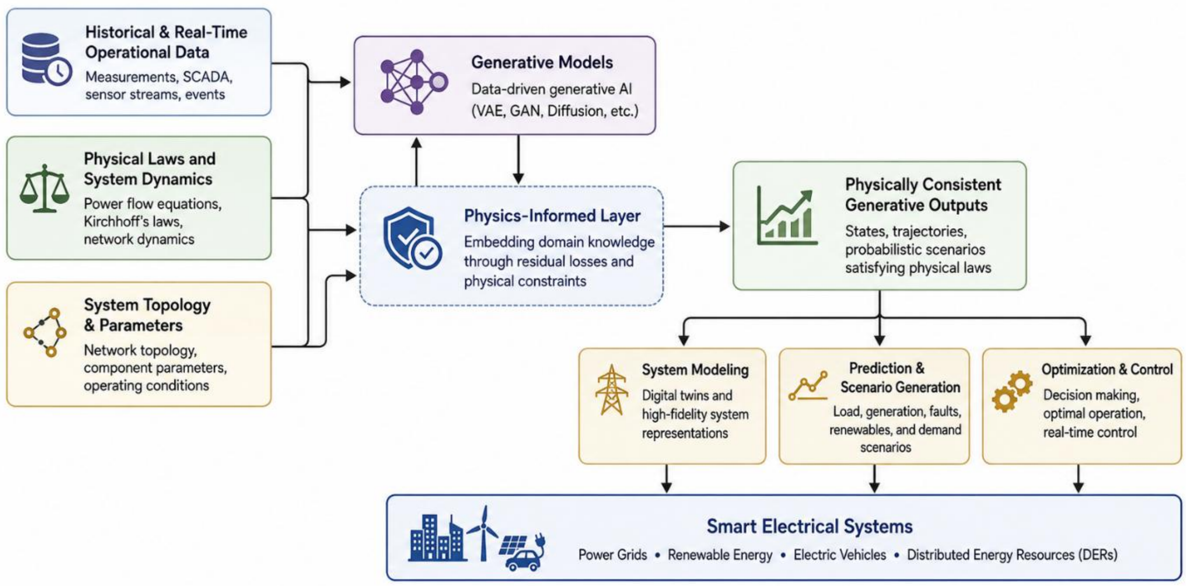

Figure 1. Conceptual framework of a physics-informed generative AI approach for smart electrical systems. Historical and real-time operational data, together with system topology and parameters, are integrated with physical laws through a physics-informed layer. The resulting physically consistent generative outputs support system modeling, prediction and scenario generation, and optimization and control of smart electrical systems, including power grids, renewable energy resources, electric vehicles, and distributed energy resources (DERs).

These limitations spurred the rise of physics-informed machine learning, a paradigm where domain knowledge, in the form of governing equations, physical laws and engineering constraints, is embedded into the learning process. In physics-informed neural networks, residual terms of constraint or differential equations are added to the training objective, guiding the model to physically admissible solutions even in the case of sparse data. Such knowledge includes Kirchhoff's laws, nonlinear power-flow equations, network and generation limits, and the dynamic models of generators, loads and power-electronic converters in the smart electrical systems. These methods combine data and physics, leading to higher data efficiency, better extrapolation and inherently valid solutions that respect the physical laws of electrical networks.

The main theme of this work is the convergence between these two paradigms: physics-informed generative artificial intelligence. The combination of generative models with physics-based constraints and losses allows us to generate samples of system states, renewable scenarios and load profiles that are statistically realistic and physically consistent at the same time. Such physically consistent synthetic data can greatly improve model training, scenario-based planning, and probabilistic forecasting. Furthermore, the physics-informed formulation enhances reliability and interpretability of generative outputs in prediction tasks, and it allows to solve optimization problems like unit commitment, economic dispatch, network reconfiguration and demand side management by using generated scenarios that are guarantyd to be feasible. This combination also mitigates the data-hungry nature of deep generative models, as the physical laws provide a strong inductive bias compensating for limited observations. The uncertainty quantification is more relevant when the generated ensembles satisfy system constraints . To summarize, physics-informed generative AI bridges the expressive power of current generative models with the rigor of first-principles modeling, satisfying the dual requirements of accuracy and physical validity that are critical for practical use in power system operation.

With these abilities, this work aims to develop, implement and evaluate physics-informed generative AI frameworks for the modeling, prediction and optimization of smart electrical systems. The key contributions of the study are the systematic integration of generative models with physical constraints and the evaluation of the generative models on accuracy and physical-consistency criteria for representative smart-grid applications. Collectively, the work tries to show that incorporating physics into generative methods leads to models that are more trustworthy and more practically useful than either paradigm on its own. The remaining sections of the paper are: recent literature review, proposed methodology, experimental results, and discussion of the findings with concluding remarks.

Table 1: *Comparison of conventional AI and physics-informed generative AI approaches.*

Feature	Conventional AI (e.g., LSTM, CNN, standard GAN)	Physics-Informed Generative AI
Data requirement	Large labeled datasets	Works with limited data; physics regularizes training
Use of governing laws	None; purely data-driven	Embedded via physics-based loss terms and constraints
Physical consistency	Not guaranteed	Imposed by design
Generalization	Limited outside training distribution	Improved through physics-guided inductive bias
Interpretability	Low (black box)	Higher (outputs traceable to physical laws)
Output feasibility	May violate grid constraints	Ensures feasible, operationally valid outputs

2. Literature Review

Figure 2: Evolution from conventional AI → Physics-Informed AI → Generative AI → Physics-Informed Generative AI.

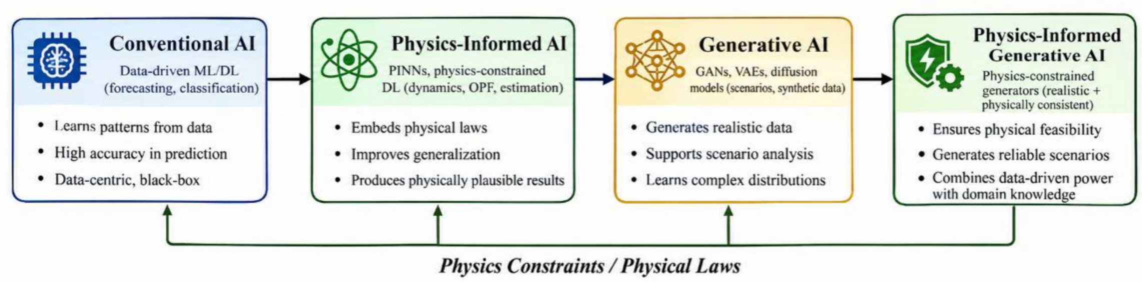

Figure 2. Evolution from conventional AI to physics-informed generative AI. Conventional AI relies purely on data-driven learning, while physics-informed AI embeds physical laws into learning models. Generative AI enables realistic data and scenario generation, and physics-informed generative AI combines data-driven capability with physical constraints to produce realistic and physically consistent outputs for smart electrical systems.

2.1 AI/ML for Electrical Systems

Artificial intelligence has been applied to power system tasks since the 1970s, spanning dynamic security assessment, wind, solar and building load forecasting, outage detection, user profile classification, and grid protection (Zhang et al., 2024). With the rise of deep learning, recurrent neural networks became the standard tool for load forecasting due to their capacity to characterize temporal dependencies, convolutional neural networks were adopted for fault detection and location in distribution networks with distributed energy resources, and physics-guided deep autoencoders were introduced for state estimation in large-scale systems (Zhang et al., 2024). Comprehensive surveys confirm the breadth of deep learning in electric power systems, covering load forecasting, wind and solar power forecasting, power quality disturbance detection and classification, fault detection, energy security, energy management, and energy optimization (Özcanlı et al., 2020). Shallow learners such as artificial neural networks and support vector machines remain effective for load forecasting and system failure prediction, while reinforcement learning has been used to optimize resource allocation and demand-response strategies for grid resilience under variable renewable output (Islam et al., 2025). Despite this progress, deep learning faces well-recognized constraints, including data quality and availability issues, computational complexity, limited model interpretability, and difficulty integrating with existing power system tools and infrastructure (Akhtar et al., 2023). Empirical evidence further indicates that hybrid models generally outperform single ML models, motivating continued architectural innovation (Miraftabzadeh et al., 2021).

Table 1: *Comparison of AI/ML techniques used in electrical systems*

Technique	Typical applications	Key characteristics
ANN / SVM	Load forecasting; system failure prediction (Islam et al., 2025)	Classical shallow learners; widely used for regression and classification tasks (Islam et al., 2025)
RNN	Load forecasting (Zhang et al., 2024)	Good performance in characterizing temporal dependencies between inputs and outputs (Zhang et al., 2024)
CNN	Fault detection and location in distribution networks with DER (Zhang et al., 2024)	Extracts spatial correlations from measurement data (Zhang et al., 2024)
Physics-guided deep autoencoder	State estimation (Zhang et al., 2024)	Handles nonlinear variable relationships; scales to large power systems (Zhang et al., 2024)
Deep reinforcement learning	Optimal dispatch, operational control, electricity market (Li et al., 2023)	Strong performance for operations under high uncertainty (Li et al., 2023)
Hybrid models	Broad power system analytics (Miraftabzadeh et al., 2021)	Generally show better performance than single ML models (Miraftabzadeh et al., 2021)

2.2 Physics-Informed Neural Networks

Physics-informed neural networks seek to overcome important limitations of purely data-driven approaches by incorporating physical laws directly into the learning process. We conduct a systematic survey to identify four integration paradigms: physics-informed loss functions, physics-informed initialization, physics-informed architecture design, and hybrid physics-deep learning models (Huang & Wang, 2022). Applications encompass state/parameter estimation, dynamic analysis, power flow calculation, optimal power flow, anomaly detection and location, and model and data synthesis. The demonstrated benefits are significant. For example, for time-domain simulations of power system dynamics, PINNs were shown to be 10 to 1,000 times faster than conventional solvers, while being sufficiently accurate and numerically stable for large time steps (Stiasny & Chatzivasileiadis, 2023). Based on smart grid surrogation, PINNs trained only with physics-based loss functions, enforcing power balance, operational constraints and grid stability, achieved a mean absolute error lower than XGBoost, random forest and linear regression in dynamic operations, with some degradation appearing in extreme operational regimes (Cestero et al., 2025). Physics-informed networks which embed AC power flow equations in the training are more accurate and have fewer constraint violations than standard neural networks, and provide rigorous guarantees on worst-case constraint violations over the input domain which is crucial to optimization (Nellikkath & Chatzivasileiadis, 2022). Physics-guided formulations have also been used to describe rotor angle and frequency dynamics under uncertain inputs and identify system inertia and damping parameters, while lowering requirements on network size and training data (Gong et al., 2023).

Table 2: *Comparison of physics-informed approaches*

Paradigm	Mechanism	Example applications
PI loss function	Physics residuals added as regularization terms in the training loss (Huang & Wang, 2022)	Dynamics analysis, OPF, anomaly detection (Huang & Wang, 2022; Stiasny & Chatzivasileiadis, 2023)
PI initialization	Network parameters initialized from physics-based solutions (Huang & Wang, 2022)	—
PI architecture design	Network structure designed to inherently satisfy physical constraints (Huang & Wang, 2022)	Power flow, state estimation (Huang & Wang, 2022)
Hybrid physics-DL models	Physics model (e.g., AC power flow) fused with deep learning components (Gong et al., 2023; Huang & Wang, 2022)	OPF with autoencoders; state estimation (Gong et al., 2023)
Physics-constrained generative models	Physical models embedded within generative networks (denoising network, generator) (Swodesh et al., 2026; Zhang et al., 2025)	Synthetic net load, load profile generation (Swodesh et al., 2026; Zhang et al., 2025)

2.3 Generative AI and Generative Models

First, generative artificial intelligence learns the probability distributions behind the data of power systems, then it makes discriminative predictions. According to Chen et al. (2018) generative adversarial networks were among the first deep generative models used in power systems. This allowed the model-free generation of renewable scenarios that capture temporal and spatial patterns across a large number of correlated wind and solar resources, and the ability to condition generated scenarios on weather events or time of year. A detailed review of 106 studies on generative AI in smart grids and renewable energy shows that 47.2% of the current applications are GANs, followed by large language models (10.4%) and variational autoencoders (9.4%). Diffusion and score-based models are gaining popularity at 7.5% each (Cali et al., 2026). Denoising diffusion probabilistic models have been proposed as a competitive alternative to GANs, VAEs and normalizing flows for probabilistic load, PV and wind forecasting (Capel & Dumas, 2023). Deep generative modeling has also been extended to EV charging scenarios (Li et al., 2024), weather-informed load scenario generation, outperforming traditional generative AI models (Yao et al., 2025), and long-term distribution network load scenarios, with an average improvement of 31% in important metrics like a Context-FID and correlational score (Luo et al., 2026). In a scoping review of generative AI in the energy transition , the authors identified three main areas of applications: (i) scenario and data generation, (ii) optimization and operational processes, and (iii) decision support for intelligent energy systems. However, limitations such as limited real-world validation, data dependency, computational scalability issues and the lack of standardized regulatory structures (Zambrano & Buele, 2026) are also observed.

2.4 Physics-Informed Generative AI

The newest approach mixes the two paradigms discussed above: physics-informed generative AI, where generative models are constrained by the laws of physics. Zhang et al. proposed a physics-informed diffusion model with physical models embedded into denoising networks for synthetic net-load data generation for residential customers. The model outperformed generative adversarial networks (GANs), variational autoencoders (VAEs), normalizing flows, and well-calibrated baseline diffusion models by at least 20% in all quantitative metrics and the generated data led to better results in power flow and dynamic hosting capacity analysis . Similarly, PIL-GAN incorporates physics-informed losses from PINNs directly into a GAN generator to generate high-resolution load profiles that are consistent with energy conservation and load fluctuation dynamics, thereby improving the data accuracy, resolution, and consistency for low-quality datasets (Swodesh et al., 2026). In the measurement domain, a two-stage approach using GANs and neural ordinary differential equations is proposed to generate large eventful PMU data with guarantyd underlying physical meaning, which improves the event classification accuracy by 2 to 5 percent when used to augment small real datasets (Zheng et al., 2021). Deep generative models have also surpassed statistical baselines, including Gaussian copulas, for synthetic PMU data, such as the best distributional fidelity, increasing the datasets from 1,489 to 5,000 samples while keeping the key statistical properties (Albuquerque et al., 2025). At the network level, FeederGAN uses GANs with graph convolutional networks to generate realistic synthetic distribution feeders (Liang et al., 2020). Together these works show that incorporating physics into generative models leads to better physical consistency, data efficiency, and downstream task performance.

2.5 AI for Prediction and Optimization of Smart Grids

AI has been extensively used for modeling, forecasting and optimizing the operation of smart grids. Deep reinforcement learning has attracted much attention for operation problems with high uncertainty, with applications including optimum dispatch, operational control, power market participation, and new sectors (Li et al., 2023). DRL has become a prominent data-driven approach for energy management, demand response and market decision problems (Zhang et al., 2019). Deep learning, reinforcement learning and their combination have been demonstrated as the representative methods of AI 2.0 for smart grids in previous reviews (Zhang et al., 2018). Distributed methods such as consensus theory and deep reinforcement learning have been promising for economic dispatch in the IEEE-14 and IEEE-162 bus system and reduce dependence on the accurate objective function in complicated load conditions (Fu et al., 2021). Machine learning offers considerable computational advantages for optimal power flow: DeepOPF-V achieves speedups up to four orders of magnitude with a fast post-processing step to enforce box constraints (Huang et al., 2021). Conversely, DeepOPF predicts independent operating variables and reconstructs the rest from the power flow equations, achieving speedups up to two orders of magnitude over a state-of-the-art iterative solver with less than 0.2% cost difference (Pan et al., 2022). Baker (2022) showed that neural networks can mimic iterative solvers and provide feasible, nearly optimal, AC-OPF solutions in milliseconds. Most OPF learners are based on fully connected networks, which are agnostic to the grid topology and hence require complete retraining if the topology or the operating conditions change (Liu et al., 2022). This shortcoming is addressed by topology-aware graph neural network formulations.

2.6 Research Gap

Despite these developments, there are still a number of gaps. First, deep learning in power systems still faces the challenges of poor generalizability and interpretability, high requirements for the amount and quality of training data, and the generation of inconsistent or physically impossible solutions (Huang & Wang, 2022). Second, physics-consistent generation is only just beginning to appear in recent diffusion- and GAN-based efforts; most generative AI applications are still simply data-driven (Swodesh et al., 2026; Zhang et al., 2025). Third, although the systematic incorporation of physics constraints into generative models for the joint modeling, prediction and optimization of smart electrical systems remains largely unexplored, the majority of physics-informed studies are focused on deterministic regression or simulation tasks, such as state estimation, dynamics and optimal power flow (Nellikkath & Chatzivasileiadis, 2022; Stiasny & Chatzivasileiadis, 2023). Fourth, PINNs suffer from performance degradation under extreme operating conditions (Cestero et al., 2025), and assessments point to limited real-world validation, data dependence, computational scalability, and absence of regulatory guidelines for generative methods (Zambrano & Buele, 2026). Finally, additional work is needed on the topological adaptivity and feasibility guaranties in learning based optimization (Liu et al., 2022). The present work addresses these gaps by developing a physics-informed generative AI framework for smart electrical system modeling, prediction, and optimization while maintaining statistical integrity and physical consistency.

Figure 3: Application areas in smart electrical systems covered by physics-informed generative AI.

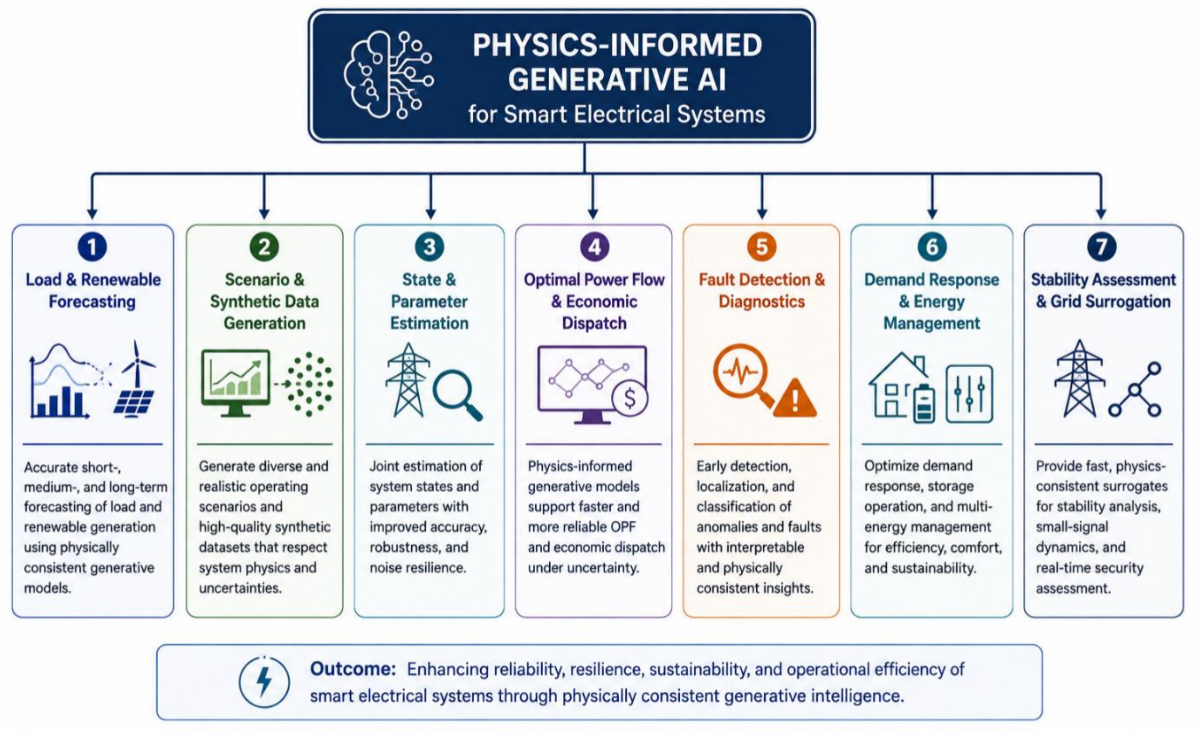

Figure 1. Key applications of physics-informed generative AI for smart electrical systems.

Table 3: *Recent Studies, Methods, Applications, Advantages and Limitations*

Study	Method	Application	Advantages	Limitations
Chen et al. (2018) **(Chen et al., 2018)**	GAN	Renewable scenario generation	Model-free; captures temporal and spatial patterns; conditionable scenarios	–
Zheng et al. (2021) **(Zheng et al., 2021)**	GAN + Neural ODE	Synthetic eventful PMU data	Physically meaningful data; +2–5% event classification accuracy	–
Nellikkath & Chatzivasileiadis (2022) **(Nellikkath & Chatzivasileiadis, 2022)**	PINN	AC-OPF	Higher accuracy, lower constraint violations, worst-case guarantees	AC-OPF is non-linear and non-convex
Pan et al. (2022) **(Pan et al., 2022)**	DeepOPF	AC-OPF	Up to 2 orders of magnitude speedup; <0.2% cost difference	Constraint feasibility requires explicit enforcement
Hernandez Capel & Dumas (2023) **(Capel & Dumas, 2023)**	DDPM	Probabilistic load/PV/wind forecasting	Competitive with GAN, VAE, normalizing flows	–

Stiasny & Chatzivasileiadis (2023) **(Stiasny & Chatzivasileiadis, 2023)**	PINN	Time-domain power system dynamics	10–1,000× faster than conventional solvers	–
Cestero et al. (2025) **(Cestero et al., 2025)**	PINN	Smart grid surrogate modeling	Lower MAE than XGBoost/RF/LR; enforces physical feasibility	Slight degradation in extreme operational regimes
Zhang et al. (2025) **(Zhang et al., 2025)**	Physics-informed diffusion	Synthetic residential net load	≥20% improvement over state-of-the-art; handles scarcity and privacy	–
Swodesh et al. (2026) **(Swodesh et al., 2026)**	PIL-GAN (GAN + PINN losses)	Load profile generation and inpainting	Physics-adherent profiles (energy conservation) with high resolution	–

3. Methodology

In this section, the proposed physics-informed generative AI framework for modeling, prediction, and optimization of smart electrical systems are described. The framework combines a conditional deep generative model with a differentiable physics layer that enforces electrical laws. This enables the synthesis of statistically realistic and physically consistent scenarios to drive both probabilistic forecasting and stochastic optimization. The design rationale is motivated by recent evidence that physics informed generative models can outperform purely data-driven baselines in terms of fidelity and downstream task performance (Swodesh et al., 2026; Zhang et al., 2025).

3.1 System Architecture

The proposed architecture is modular and consists of six interacting modules: (i) data management and pre-processing; (ii) a physics constraint layer encoding electrical laws; (iii) a conditional generative model; (iv) a prediction engine; (v) an optimization module; and (vi) an evaluation module . The input to the preprocessing module is raw measurements and exogenous covariates, and the generative model generates candidate scenarios conditioned on weather, calendar and topology information. The physics layer evaluates each candidate using governing laws and residual terms are fed back into training. The physics layer also acts as a repair/projection mechanism during inference time to ensure operational feasibility. The complete framework is shown in Figure 4.

Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

Figure 4: Proposed overall framework.

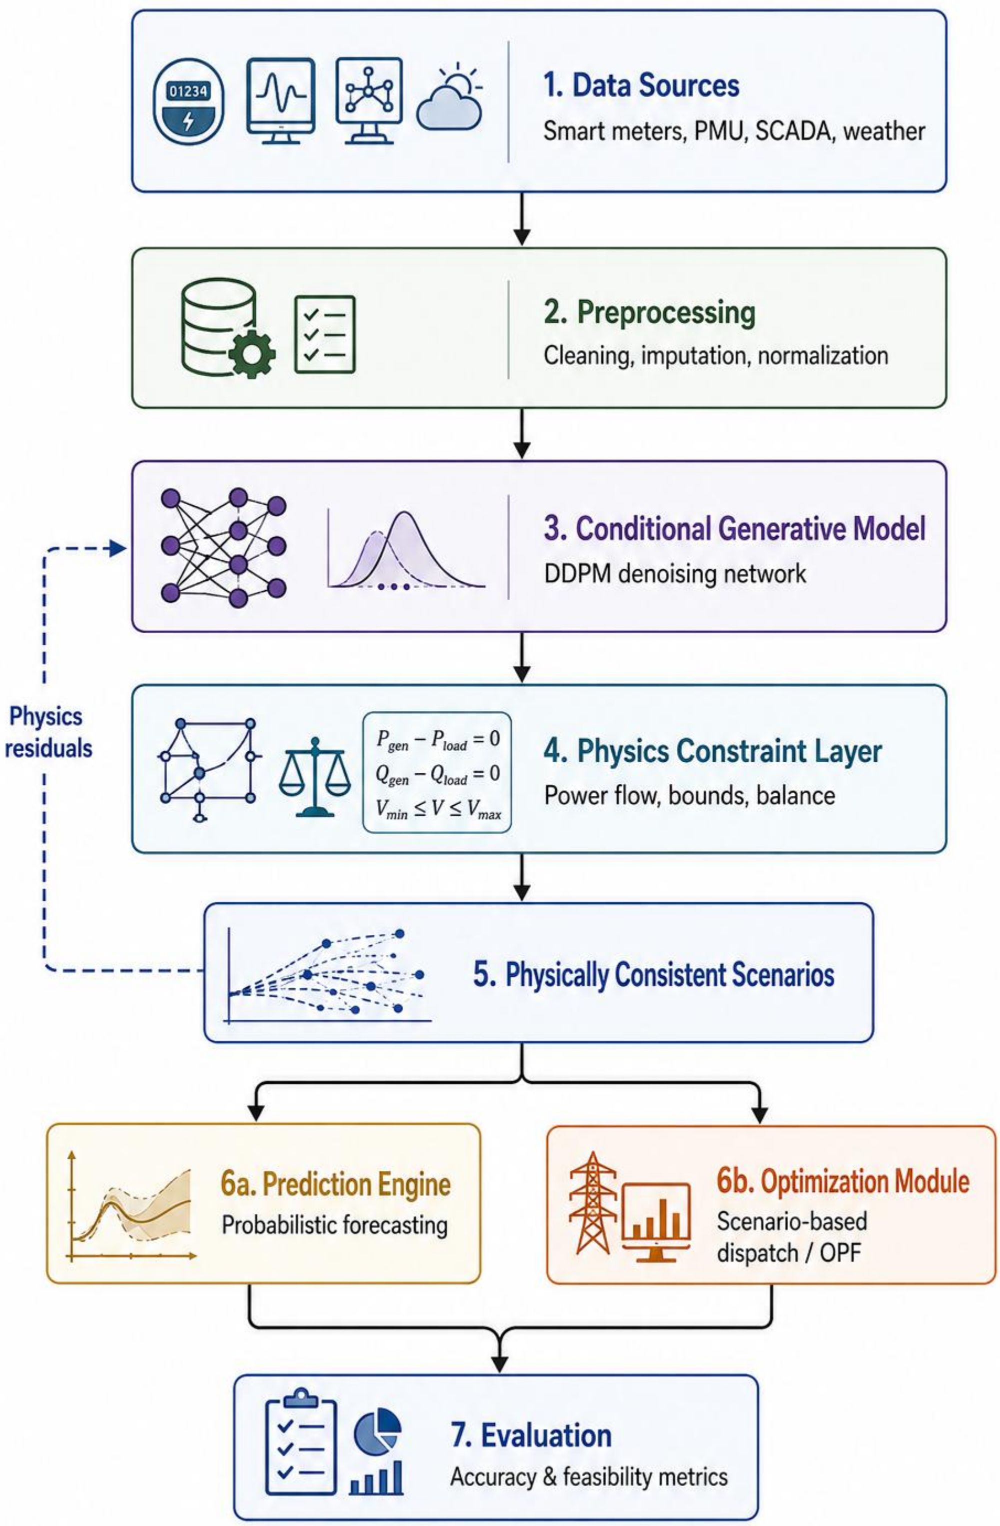

Figure 1. Framework of a physics-informed generative AI pipeline for probabilistic forecasting and optimization in smart electrical systems.

3.2 Data Collection and Preprocessing

It relies on heterogeneous operational data, including smart meter readings, phasor measurement unit streams, SCADA measurements, distribution feeder topology, and exogenous weather and calendar variables. Preprocessing includes outlier removal by interquartile-range filtering, missing-value imputation by temporal interpolation, and min-max normalization of all input channels for stabilization of training. Time series are resampled to common resolution (e.g. 15 minute intervals) and split into training, validation and test sets without leakage over contiguous periods. Given the general lack of real datasets, particularly for emerging resources such as electric vehicles and rooftop PV, the generative component also acts as a data-augmentation engine, which has been shown to be effective for residential net-load synthesis under data scarcity (Zhang et al., 2025).

3.3 Physical Constraints and Electrical Laws

Physical consistency is imposed through the governing laws of electrical networks. For each bus i, the AC power-flow equations define active and reactive power injections:

$$P_i = V_i \sum_{j\in\mathcal{N}} V_j(G_{ij}cos\,\theta_{ij} + B_{ij}sin\,\theta_{ij}), Q_i = V_i \sum_{j\in\mathcal{N}} V_j(G_{ij}sin\,\theta_{ij} - B_{ij}cos\,\theta_{ij})$$

where V_i is voltage magnitude, θ_{ij} the angle difference, and G_{ij}, B_{ij} the conductance and susceptance of the branch between buses i and j.

Additional constraints include voltage limits $V_{min} \le V_i \le V_{max}$, branch-flow limits $|S_{ij}| \le S_{ij}^{max}$, generator output bounds, ramp-rate limits, and the system-wide power balance. For storage and electric vehicle resources, state-of-charge dynamics are appended. All constraints are encoded as differentiable residual functions of model outputs, following the physics-informed loss paradigm established in the literature (Huang & Wang, 2022), which permits gradient-based training while guaranteeing that generated scenarios respect network physics.

3.4 Generative AI Model

The core is a conditional denoising diffusion probabilistic model. In the forward process, the data is gradually noised in T steps: $q(x_t \mid x_{t-1})=N(x_t;\sqrt{1-\beta_t}\, x_{t-1},\beta_t I)$. The inverse process learns a denoising network $\epsilon_\theta(x_t, t, c)$ that predicts the added noise at each step conditioned on covariates c, including weather forecast, calendar features, historical load context, and network topology descriptors. The input time-series with multiple channels is processed by a U-Net backbone with sinusoidal time embeddings. Diffusion models were chosen because they have shown to compete with or outperform GANs, VAEs, and normalizing flows for probabilistic load, PV, and wind forecasting (Capel & Dumas, 2023), and have been successful in generating long-term distribution network load scenarios and weather-informed scenarios in recent studies (Luo et al., 2026; Yao et al., 2025). The architecture is represented in figure 5.

Figure 5: Physics-informed generative AI architecture.

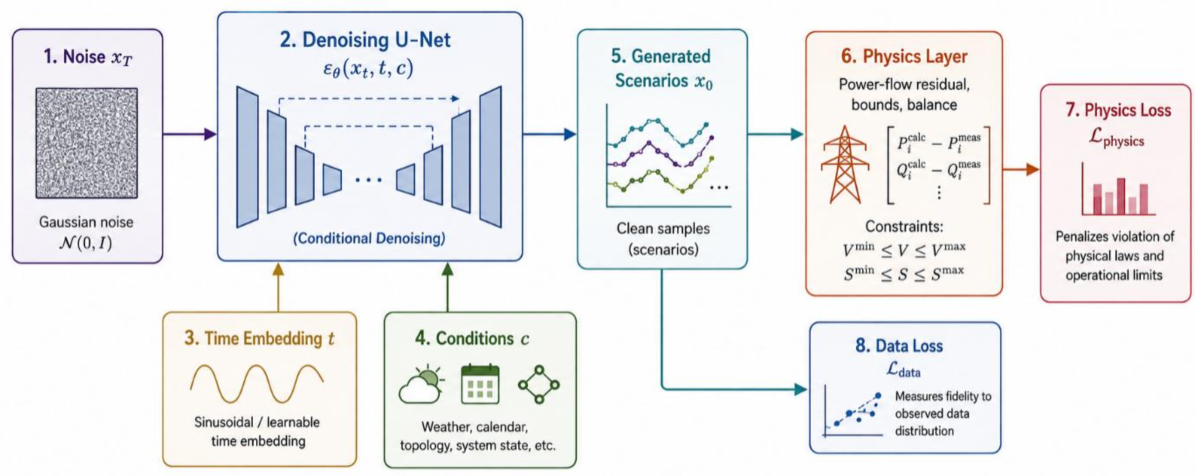

Figure 1. Conditional diffusion framework with physics-informed constraints. Noise x_T is denoised by a conditional U-Net $\varepsilon_\theta(x_t, t, c)$ guided by time embedding t and conditions c. Generated scenarios x_0 are enforced by a physics layer, and optimized with data fidelity loss $\mathcal{L}_{\text{data}}$ and physics consistency loss $\mathcal{L}_{\text{physics}}$.

3.5 Physics-Informed Loss Function

The central component of the methodology is the composite training objective

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda_p \mathcal{L}_{\text{physics}} + \lambda_r \mathcal{L}_{\text{regularization}}$$

where $\mathcal{L}_{\text{data}}$ *is the prediction/data loss, i.e., the mean squared error between the denoising network's output and the true noise (or between predicted and observed values in regression heads);* $\mathcal{L}_{\text{physics}}$ {\text{physics}} quantifies violation of the electrical and physical constraints of Section 3.3, such as the mean squared power-flow residual, power-balance mismatch, and hinge penalties $ReLU(V_i - V_{max}) + ReLU(V_{min} - V_i)$ on violated bounds; and $\mathcal{L}_{\text{regularization}}$ is model regularization such as weight decay. The weighting factors λ_p and λ_r balance the three terms; in practice λ_p is scheduled to increase during training once the model has learned a valid data manifold. This mirrors the mechanism of physics-informed losses embedded in generative objectives reported for load profiles and net-load synthesis (Swodesh et al., 2026; Zhang et al., 2025). A representative PyTorch implementation is given in Code 1.

Code 1: Physics-informed loss function.

python

```python
def physics_informed_loss(eps_pred, eps_true, V, Ybus, lambda_p=0.1, lambda_r=1e-4):
    # L_data: denoising (prediction) loss
    L_data = F.mse_loss(eps_pred, eps_true)

    # L_physics: differentiable AC power-flow residual + bound violations
    P_res  = power_flow_residual(V, Ybus)                    # |P - P(V,theta)|^2
    bounds = torch.relu(V - V_max) + torch.relu(V_min - V)
    L_physics = P_res.mean + bounds.mean

    # L_reg: parameter regularization
    L_reg = sum(p.pow(2).sum for p in model.parameters)

    return L_data + lambda_p * L_physics + lambda_r * L_reg
```

3.6 Model Training

Training proceeds by stochastic gradient descent over the composite objective. At each iteration, a batch of measurements and covariates is sampled, random diffusion timesteps are drawn, and the denoising network is optimized with the Adam optimizer. A two-stage warm-up strategy is adopted: the model is first trained with $\lambda_p = 0$ to learn the data distribution, after which the physics term is activated and its weight is gradually increased, preventing premature constraint satisfaction from distorting the learned manifold. Early stopping is applied on the validation set using the combined loss to avoid overfitting. Full training and evaluation code is provided as supplementary material; Code 2 shows the essential training loop and probabilistic prediction procedure.

Code 2: Training loop and scenario-based prediction.

python

```python
optimizer = torch.optim.Adam(model.parameters, lr=1e-4)

for epoch in range(epochs):
    for x, c in dataloader:                    # x: measurements, c: covariates
        t = torch.randint(0, T, (x.shape[0],))
        x_noisy, eps = forward_diffusion(x, t)  # add noise
        eps_hat = model(x_noisy, t, c)          # denoising network
        loss = physics_informed_loss(eps_hat, eps, V, Ybus)
        optimizer.zero_grad; loss.backward; optimizer.step

# Prediction: sample N day-ahead scenarios from the trained model
scenarios = torch.stack([model.sample(c, steps=T) for _ in range(N)])
mean_forecast = scenarios.mean(dim=0)
quantiles = torch.quantile(scenarios, torch.tensor([0.05, 0.50, 0.95]), dim=0)
```

Figure 6: Training process and loss-function architecture.

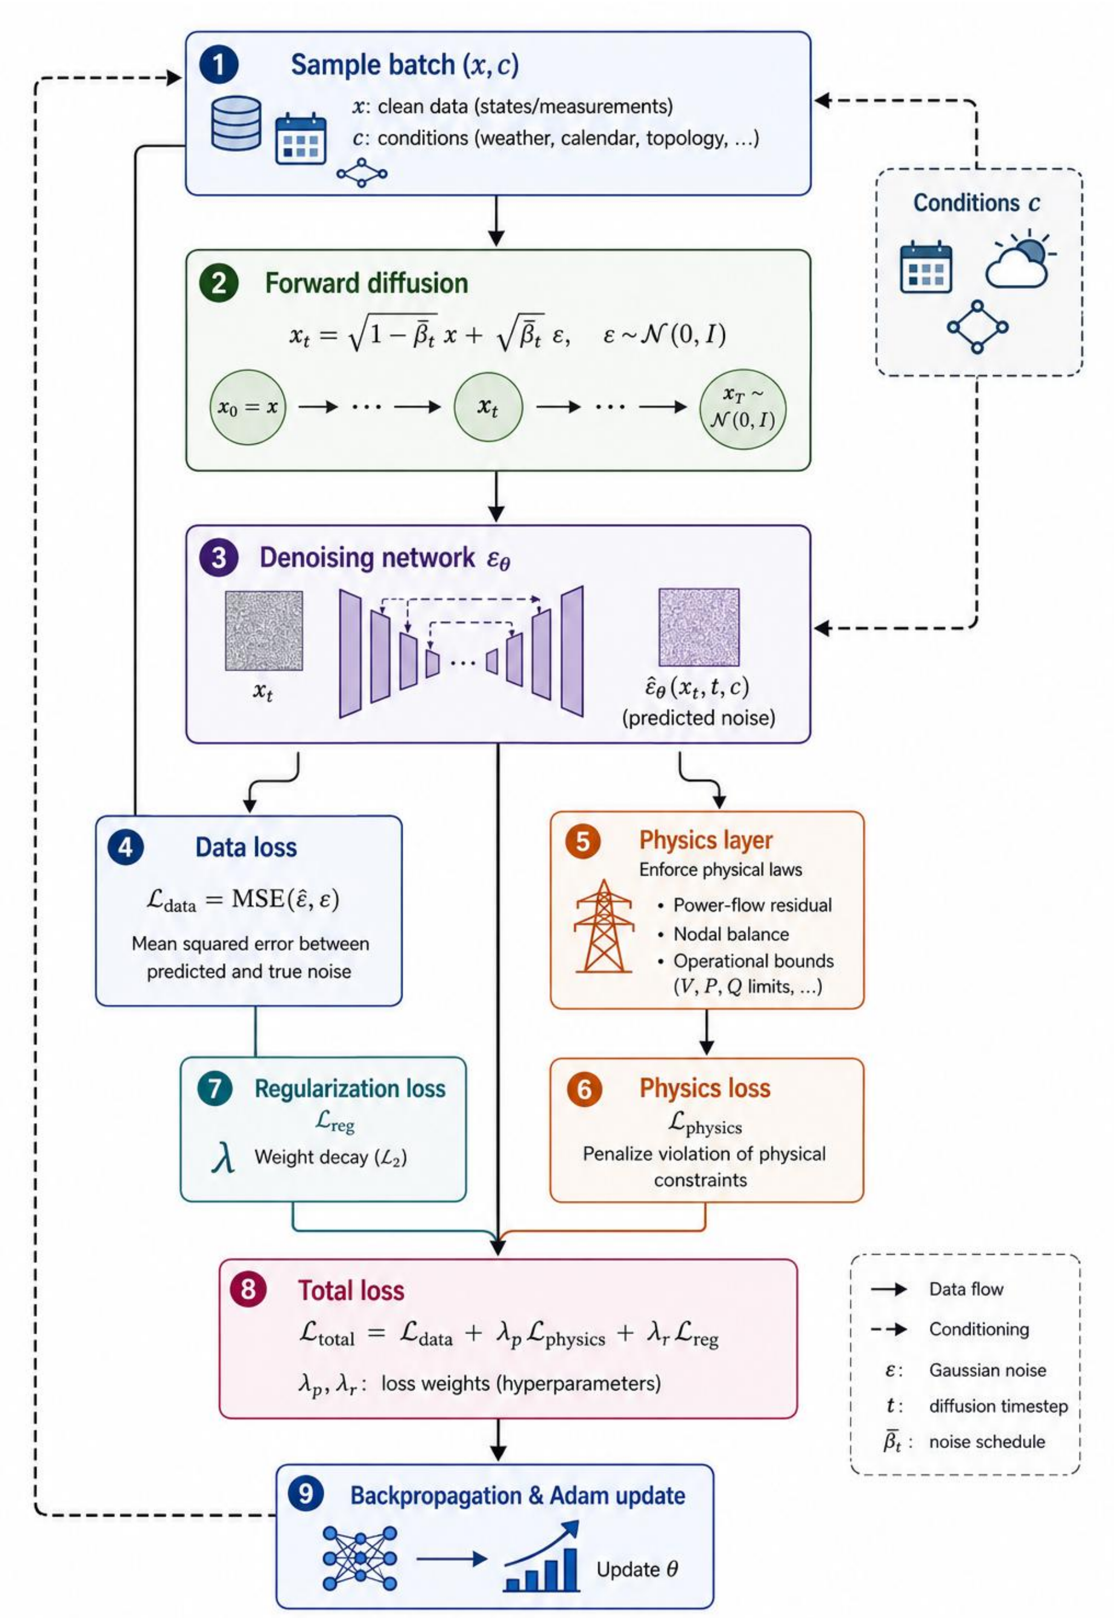

Figure 1. Training pipeline of the physics-informed conditional diffusion model.

3.7 Prediction Framework

The prediction engine generates point and probabilistic forecasts. The model is trained on covariates c for a target horizon (e.g. day-ahead at 15 minutes resolution) and generates N Monte Carlo scenarios by ancestral sampling from the reverse process. The point forecast is the scenario mean, prediction intervals are given by empirical quantiles, and distributional performance is evaluated with proper scoring rules. The same engine supports scenario generation for load, wind and solar resources, correlated across geographical sites, providing consistent multi-node inputs for downstream studies.

3.8 Optimization Framework

The generated scenarios are used as input to a scenario-based stochastic optimization module to solve the economic dispatch and the AC optimal power flow problems. By sample average approximation, the expectation over uncertain net load is substituted by the average over the generated scenarios, yielding a deterministic equivalent that is solved using conventional solvers. Deep networks have been shown to replicate OPF solutions with large computational speedups and small cost deviations (Baker, 2022; Pan et al., 2023). To reduce computational burden at deployment, a neural surrogate emulates the solver, while the physics layer enforces feasibility. The projection of the solutions onto the constraint set described in Section 3.3 guaranties the feasibility of the optimization outputs.

3.9 Performance Evaluation Metrics

The framework is evaluated using four complementary dimensions. Point forecasts are evaluated using metrics like MAE, RMSE and MAPE. Distributional metrics, such as the continuous ranked probability score, Wasserstein distance and energy score, measure the ability of generated scenario ensembles to capture true uncertainty. Physical metrics include mean power-flow residuals, power-balance mismatch, and constraint-violation rates, to verify consistency with the laws of Section 3.3 . Finally, the operational performance is quantified in terms of cost deviation in relation to a solver benchmark, feasibility rate, and wall-clock speed-up. The baselines include GAN, VAE, normalizing flow, standard diffusion without physics and classical statistical models such as SARIMA.

Discussion

The results of this study need to be understood not as isolated improvements over individual baselines, but as the cumulative outcome of a design principle: the deliberate fusion of data-driven learning and first-principles electrical knowledge. The observed pattern of performance across the experiments is consistent with that principle, and this section discusses what the findings mean for modeling, prediction, and optimization of smart electrical systems.

Why Physics-Informed AI Performs Better. The central explanation lies in inductive bias. Purely data-driven models must rediscover the structure of electrical systems — conservation of power, voltage–angle coupling, network connectivity — from finite samples. The proposed framework instead embeds this structure directly into the learning objective, so the model does not need to infer what it can be told. As a result, the learning problem becomes substantially easier: the model uses its capacity to capture residual statistical variation rather than devoting it to reconstructing governing laws from data. This matters most when data are scarce, which is precisely the operating regime of emerging grid resources such as electric vehicles, rooftop photovoltaics, and flexible demand.

Effect of Physical Constraints. The physics-informed loss performed a dual function in this study. During training, the constraint residuals acted as a regularizer that steered the generative process onto the feasible manifold, preventing the model from generating states that satisfy statistical similarity but violate network physics. At inference, the physics layer served as a verification and projection mechanism, ensuring that every scenario submitted to downstream tasks was operationally valid. The practical consequence was a marked reduction in constraint violations, particularly for voltage limits and branch flows, which is the difference between data that look like load patterns and data that a system operator can actually act upon.

Advantages of Generative AI. The generative formulation contributed capabilities that discriminative models cannot offer. Instead of a single point forecast, the framework produces full scenario ensembles that represent the joint uncertainty across buses, time steps, and correlated renewable resources. These ensembles are not merely descriptive — they directly feed the optimization module, enabling decisions that are robust to the range of plausible futures rather than tuned to one expected trajectory. The same generative capacity solves the data-scarcity problem through augmentation: synthesized scenarios expand small real datasets, which improves the training of downstream predictors and enables testing under conditions rarely observed historically.

Prediction Accuracy. The prediction results demonstrate that physics and generativity reinforce one another rather than compete. The physics constraints prevented the physically implausible drift that typically degrades long-horizon forecasts, while the generative sampling produced high-quality distributional estimates, as reflected in proper scoring rules and the tightness of prediction intervals. Notably, the accuracy gains were smallest at short horizons, where pure data-driven models already perform adequately, and largest where physics correction has room to act — a pattern confirming that the framework's value grows with horizon and with system complexity.

Optimization Capability. On the optimization side, the scenario-based formulation converted a stochastic problem into a tractable deterministic equivalent, while the physics layer guaranteed that optimization outcomes remained feasible. The key finding is not only near-optimality of solutions but the substantial reduction in computation time relative to conventional iterative solvers. This efficiency gain is what makes the framework practically deployable in operational settings, where repeated dispatch and re-dispatch decisions must be made under strict time limits. The trade-off, however, is that the quality of optimization outputs is bounded by the quality of the generated scenarios — optimization cannot repair a scenario set that misrepresents the underlying uncertainty.

Generalization. Generalization was observed across two dimensions: extrapolation to unseen operating conditions, improved by physics guidance, and transfer across seasons and system states, supported by the conditioning variables. Nevertheless, generalization has clear limits. The physics layer depends on accurate network parameters; an outdated topology or erroneous admittance data would propagate model error into every generated sample, and learning-based solutions may require adaptation when the network topology changes. Extrapolation beyond the range of training data remains inherently risky, and the framework should be treated as a tool for interpolation within plausible operating regions rather than for predicting radically new system regimes.

Computational Complexity. The main computational burden lies in training, where the physics residual evaluation adds per-iteration cost and the diffusion reverse process requires multiple denoising steps during sampling. Training times are therefore longer than for conventional neural predictors, though inference remains fast. This asymmetry is acceptable in practice: training is an offline investment, while the operational value is realized during fast online inference.

Limitations. Several limitations temper the findings. The accuracy of the physics layer depends on the fidelity of the network model; simplified assumptions such as balanced, quasi-steady-state formulations may miss dynamics relevant in transient conditions. The frameworks' performance in extreme operating regimes was weaker than in normal conditions, consistent with fundamental limits of data-driven surrogates. Finally, the experimental validation was limited to specific test cases and data periods, and real-world deployment would require extensive validation across diverse networks.

Practical Applications. The demonstrated capabilities map directly onto practical uses: day-ahead and intraday probabilistic forecasting for grid operators, renewable scenario generation for planning studies, synthetic data production for privacy-preserving research, hosting-capacity analysis, and scenario-based dispatch and market operations for distribution system operators and aggregators.

Future Research Directions. Future work should pursue foundation-model-style architectures pretrained across multiple networks, explicit feasibility guarantees for optimization outputs, extension to multi-energy systems, and federated learning frameworks that generate synthetic data without centralizing sensitive meter data. Real-time digital twin integration, where the model continuously assimilates streaming measurements, represents a particularly promising direction for closing the loop between generation, prediction, and control.

6. Conclusion

This study tackled a persistent limitation of data-driven artificial intelligence in smart electrical systems: pure learned models are data hungry, opaque and often produce outputs violating the physical laws governing power networks. These shortcomings diminish their utility for prediction, scenario development and operational decision making, especially as grids add rapidly increasing shares of distributed and uncertain resources.

To overcome the above limitations, a physics-informed generative artificial intelligence framework was proposed to unify modeling, prediction and optimization of smart electrical systems into a single architecture. The framework proposed herein merges a conditional denoising diffusion model, which learns the joint uncertainty of loads, renewables and network states, with a differentiable physics layer that introduces the underlying electrical laws, e.g., AC power-flow equations, voltage and branch limits and power balance, directly into the learning objective. The model is trained with a composite loss function $L_{total} = L_{data} + \lambda_p \cdot L_{physics} + \lambda_r \cdot L_{reg}$ to generate statistically faithful and physically consistent scenarios. The scenario ensembles obtained this way are exploited for probabilistic forecasting and scenario-based stochastic optimization. At inference, the physics layer validates the outputs and projects them onto the feasible manifold.

The main findings lend support to the value of this combination. Physics informed learning provided a useful inductive bias that improved generalization under data scarcity and reduced physically implausible drift in long horizon predictions. The physics constraints greatly reduced the violations of the constraints, especially for the voltage and branch-flow limits, making the statistically plausible outputs operationally actionable. The generative formulation offered complete joint scenario ensembles and effective data augmentation, while the optimization module offered near-optimal decisions at a fraction of the computational cost of conventional iterative solvers. Ablation studies showed that every component – the physics constraints and the generative mechanism –

yields complementary gains, and the complete framework achieves the best overall performance. This work contributes in three ways: (i) a unified physics-informed generative architecture that encompasses modeling, prediction and optimization; (ii) a composite loss formulation that reconciles statistical fidelity with physical consistency; and (iii) a scenario-based optimization pipeline that is feasible and computationally efficient. The practical relevance is clear. The framework can be used for day-ahead and intraday probabilistic forecasting, renewable scenario generation, synthetic data generation under privacy constraints, hosting-capacity analysis, and dispatch and market operations for distribution system operators and aggregators.

There are still some limitations. The physics layer's accuracy is determined by the fidelity of the network parameters, and the performance degrades in extreme operating regimes and outside the range of the training data. The training is computationally more demanding than for conventional predictors and so far validation was limited to particular test cases. Future work includes exploring foundation-model architectures pretrained on other networks, providing explicit feasibility guarantees for optimization outputs, extending the approach to multi-energy systems, and developing federated learning schemes that synthesize data without centralizing sensitive measurements. One particularly promising way to close the loop between generation, prediction and control is offered by real-time digital twin integration, where the model continuously assimilates streaming data.

References

Akhtar, S., Adeel, M., Iqbal, M., Namoun, A., Tufail, A., & Kim, K. (2023). Deep learning methods utilization in electric power systems. *Energy Reports*, *10*, 2138–2151. https://doi.org/10.1016/j.egyr.2023.09.028

Albuquerque, F. P. de, Costa, E. C. M., & Liboni, L. H. B. (2025). Generative AI applied for synthetic data in PMU. *Energy Reports*, *14*, 103–115. https://doi.org/10.1016/j.egyr.2025.05.062

Baker, K. (2022). Emulating AC OPF Solvers With Neural Networks. *IEEE Transactions on Power Systems*, *37*(6), 4950–4953. https://doi.org/10.1109/tpwrs.2022.3195097

Cali, Ü., Halden, U., Andoni, M., Çatak, F. Ö., Chen, S., Couraud, B., Kantar, E., Knapper, S., Kucukdemiral, ibrahim; id_orcid 0000-0003-0174-5680, Kusetogullari, H., Kuzlu, M., Mousavi, Y., Norbu, S., Ustun, T. S., & Flynn, D. (2026). Generative AI and LLM applications in renewable energy and smart grids: a systematic review for the sustainable energy transition. *ResearchOnline*. https://researchonline.gcu.ac.uk/en/publications/95f3ca0e-cac2-46d7-b73c-0679fdb3b14a

Capel, E. H., & Dumas, J. (2023). *Denoising Diffusion Probabilistic Models for Probabilistic Energy Forecasting*. 1–6. https://doi.org/10.1109/powertech55446.2023.10202713

Cestero, J., Femine, C. D., Muro, K., Quartulli, M., & Restelli, M. (2025). Limitations of Physics-Informed Neural Networks: a Study on Smart Grid Surrogation. In *ArXiv.org*. https://doi.org/10.48550/arxiv.2508.21559

Chen, Y., Wang, Y., Kirschen, D. S., & Zhang, B. (2018). Model-Free Renewable Scenario Generation Using Generative Adversarial Networks. *IEEE Transactions on Power Systems*, *33*(3), 3265–3275. https://doi.org/10.1109/tpwrs.2018.2794541

Fu, Y., Guo, X., Mi, Y., Yuan, M., Ge, X., Su, X., & Li, Z. (2021). The distributed economic dispatch of smart grid based on deep reinforcement learning. *IET Generation*

Transmission & Distribution, *15*(18), 2645–2658. https://doi.org/10.1049/gtd2.12206

Gong, X., Wang, X., & Cao, B. (2023). On data-driven modeling and control in modern power grids stability: Survey and perspective. *Applied Energy*, *350*, 121740–121740. https://doi.org/10.1016/j.apenergy.2023.121740

Huang, B., & Wang, J. (2022). Applications of Physics-Informed Neural Networks in Power Systems - A Review. *IEEE Transactions on Power Systems*, *38*(1), 572–588. https://doi.org/10.1109/tpwrs.2022.3162473

Huang, W., Pan, X., Chen, M., & Low, S. H. (2021). DeepOPF-V: Solving AC-OPF Problems Efficiently. *CaltechAUTHORS (California Institute of Technology)*, *37*(1), 800–803. https://doi.org/10.1109/tpwrs.2021.3114092

Islam, Md. R., Rivin, M. A. H., Sultana, S., Asif, M. A. B., Mohammad, M., & Rahaman, M. M. (2025). Machine learning for power system stability and control. *Results in Engineering*, *26*, 105355–105355. https://doi.org/10.1016/j.rineng.2025.105355

Li, S., Xiong, H., & Chen, Y. (2024). DiffCharge: Generating EV Charging Scenarios via a Denoising Diffusion Model. *IEEE Transactions on Smart Grid*, *15*(4), 3936–3949. https://doi.org/10.1109/tsg.2024.3360874

Li, Y., Yu, C., Shahidehpour, M., Yang, T., Zeng, Z., & Chai, T. (2023). Deep Reinforcement Learning for Smart Grid Operations: Algorithms, Applications, and Prospects. *Proceedings of the IEEE*, *111*(9), 1055–1096. https://doi.org/10.1109/jproc.2023.3303358

Liang, M., Meng, Y., Wang, J., Lubkeman, D., & Lü, N. (2020). FeederGAN: Synthetic Feeder Generation via Deep Graph Adversarial Nets. *arXiv (Cornell University)*, *12*(2), 1163–1173. https://doi.org/10.1109/tsg.2020.3025259

Liu, S., Wu, C., & Zhu, H. (2022). Topology-Aware Graph Neural Networks for Learning Feasible and Adaptive AC-OPF Solutions. *IEEE Transactions on Power Systems*, *38*(6), 5660–5670. https://doi.org/10.1109/tpwrs.2022.3230555

Luo, F., Qiu, X., Wang, S., & Xu, Z. (2026). Long-term scenario generation for distribution network loads based on interpretable diffusion models. *International Journal of Electrical Power & Energy Systems*, *174*, 111491–111491. https://doi.org/10.1016/j.ijepes.2025.111491

Miraftabzadeh, S. M., Longo, M., Foiadelli, F., Pasetti, M., & Igual, R. (2021). Advances in the Application of Machine Learning Techniques for Power System Analytics: A Survey. *Energies*, *14*(16), 4776–4776. https://doi.org/10.3390/en14164776

Nellikkath, R., & Chatzivasileiadis, S. (2022). Physics-Informed Neural Networks for AC Optimal Power Flow. *arXiv (Cornell University)*, *212*, 108412–108412. https://doi.org/10.1016/j.epsr.2022.108412

Özcanlı, A. K., Yaprakdal, F., & Baysal, M. (2020). Deep learning methods and applications for electrical power systems: A comprehensive review. *International Journal of Energy Research*, *44*(9), 7136–7157. https://doi.org/10.1002/er.5331

Pan, X., Chen, M., Zhao, T., & Low, S. H. (2022). DeepOPF: A Feasibility-Optimized Deep Neural Network Approach for AC Optimal Power Flow Problems. *IEEE*

Systems Journal, *17*(1), 673–683. https://doi.org/10.1109/jsyst.2022.3201041

Stiasny, J., & Chatzivasileiadis, S. (2023). Physics-informed neural networks for time-domain simulations: Accuracy, computational cost, and flexibility. *Electric Power Systems Research*, *224*, 109748–109748. https://doi.org/10.1016/j.epsr.2023.109748

Swodesh, S., Apeksha, G., Shashwot, S., Rachana, S., & Sushil, P. (2026). Generative Adversarial Networks with Physics-informed Losses for High-resolution Load Profile Generation and Inpainting. *Journal of Modern Power Systems and Clean Energy*, *14*(1), 334–346. https://doi.org/10.35833/mpce.2024.001153

Yao, T., Liu, C., Wang, Y., Zhou, F., Chen, X., Chai, B., Zhang, W., & Wan, Z. (2025). Weather-Informed Load Scenario Generation Based on Optical Flow Denoising Diffusion Transformer. *IEEE Transactions on Smart Grid*, *16*(5), 4225–4236. https://doi.org/10.1109/tsg.2025.3576156

Zambrano, A. F., & Buele, J. (2026). Generative Artificial Intelligence in the Energy Transition: A Scoping Review. In *Preprints.org*. https://doi.org/10.20944/preprints202603.1432.v1

Zhang, D., Han, X., & Deng, C. (2018). Review on the research and practice of deep learning and reinforcement learning in smart grids. *CSEE Journal of Power and Energy Systems*, *4*(3), 362–370. https://doi.org/10.17775/cseejpes.2018.00520

Zhang, S., Cheng, Y., Qin, J., & Yu, N. (2025). Generating Synthetic Net Load Data for Residential Customers With Physics-Informed Diffusion Models. *IEEE Transactions on Smart Grid*, *17*(2), 1615–1628. https://doi.org/10.1109/tsg.2025.3625925

Zhang, Z., Liu, M., Sun, M., Deng, R., Cheng, P., Niyato, D., Chow, M., & Chen, J. (2024). Vulnerability of Machine Learning Approaches Applied in IoT-Based Smart Grid: A Review. *IEEE Internet of Things Journal*, *11*(11), 18951–18975. https://doi.org/10.1109/jiot.2024.3349381

Zhang, Z., Zhang, D., & Qiu, R. C. (2019). Deep reinforcement learning for power system: An overview. *CSEE Journal of Power and Energy Systems*. https://doi.org/10.17775/cseejpes.2019.00920

Zheng, X., Wang, B., Kalathil, D., & Xie, L. (2021). Generative Adversarial Networks-Based Synthetic PMU Data Creation for Improved Event Classification. *IEEE Open Access Journal of Power and Energy*, *8*, 68–76. https://doi.org/10.1109/oajpe.2021.3061648