

AI-Based Digital Twins for Renewable Energy Grid Optimization: A Framework for Real-Time Forecasting, State Estimation and Adaptive Dispatch

Dr. G. Sripriya¹, V. S. Guhan², A. S. Nandha Kisore³

¹Assistant Professor, Department of Computer Applications

^{2,3}Student, Department of Computer Applications

Sri Krishna Arts and Science College, Coimbatore, India

sripriyag@skasc.ac.in, guhanvs24bca113@skasc.ac.in, nandhakisoreas24bca135@skasc.ac.in

Abstract: Growing reliance on distributed, weather-dependent renewable sources brings sizable variability, uncertainty, and stability problems to today's power grids. This work proposes an AI-driven digital twin (DT) framework that pairs a real-time physical-data grid model with a layered AI optimisation engine made up of an LSTM/Transformer forecaster, a graph neural network (GNN) state estimator, and a deep reinforcement learning (DRL) dispatch controller. The twin continually mirrors the physical grid, supporting scenario simulation, predictive fault detection, and closed-loop control recommendations before any action is applied to the real system. A case study on a modified IEEE 33-bus feeder with solar PV and wind generation shows the framework raises renewable utilisation by 19 points, cuts curtailment by 28 points, and lowers fault response time by 68.2% versus a conventional energy-management system. These results suggest AI-enhanced digital twins are a scalable route to self-optimising, resilient renewable-heavy grids.

Keywords: Digital twin, artificial intelligence, renewable energy, smart grid, deep reinforcement learning, graph neural network, load forecasting, grid optimisation.

I. INTRODUCTION

The worldwide shift to decarbonised electricity has fuelled fast growth in solar PV and wind capacity. Yet the intermittent, stochastic nature of these sources creates forecasting uncertainty, voltage swings, and frequency instability that conventional SCADA-based energy management systems (EMS) struggle to handle in real time.

Digital twin (DT) technology - a virtual replica of a physical asset kept in sync via sensor data - offers a promising fix. Paired with AI, a digital twin moves from passive monitoring to an active decision-support and autonomous-control platform capable of forecasting, what-if simulation, and closed-loop optimisation.

Utilities face growing pressure to integrate distributed energy resources, EV charging, and demand-response programmes while still meeting regulator-mandated reliability targets. An AI-enabled digital twin closes this gap by continuously learning from operational data and suggesting corrective actions before, not after, a disturbance occurs.

This paper offers four contributions: (i) a layered AI-digital-twin architecture unifying forecasting, state estimation, and dispatch optimisation; (ii) an integration method linking physical grid telemetry to the twin via a scenario-simulation sandbox that ensures AI-recommended actions are safe; (iii) a quantitative case study benchmarking the framework against a conventional EMS on a high-renewable feeder; and (iv) a discussion of deployment challenges relevant to real-world adoption.

The rest of the paper is structured as follows. Section II covers related work. Section III presents the architecture. Section IV covers methodology. Section V reports results. Section VI covers implementation challenges, and Section VII concludes.



II. RELATED WORK

A. Digital Twins in Power Systems

Early digital-twin efforts in the power sector centred on asset-level monitoring via physics-based models refreshed periodically from sensor data. Recent work extends this to network-level visualisation, but most such systems stay observational rather than autonomously recommending or executing control.

B. AI-Based Forecasting and State Estimation

Machine-learning forecasts of renewable generation using LSTM and attention-based Transformers reach high short-horizon accuracy. Graph-based learning has separately been used for topology-aware state estimation and fault localisation, leveraging network graph structure for robustness under partial observability.

C. Reinforcement Learning for Grid Dispatch

Reinforcement learning (RL) has been applied to microgrid energy management, demand response, and storage dispatch, typically framed as a Markov decision process. These studies show gains over rule-based heuristics but are usually trained apart from a live, synchronised grid representation, limiting confidence in real deployment.

The framework proposed here differs by unifying forecasting, state estimation, and dispatch optimisation inside one digital twin kept synchronised with the physical grid at all times, letting the RL policy be validated in a simulation sandbox before actuation.

D. Research Gap and Positioning

Table I compares the proposed framework with representative prior studies across four dimensions: continuous twin synchronisation, joint forecasting/state-estimation training, a safety-validation sandbox, and a quantified case study. While individual AI components are well studied, their joint integration in a twin architecture that also guarantees action safety before dispatch is comparatively unexplored - the gap this paper addresses.

TABLE I POSITIONING OF THE PROPOSED FRAMEWORK RELATIVE TO PRIOR WORK

Study	Continuous Twin Sync	Joint Forecast + State Est.	Safety Sandbox + Case Study
Grieves & Vickers [1]; Tao et al. [2]	No	No	No
Wang et al. [3] (DT review)	Partial	No	No
Wu et al. [4] (Transformer)	No	N/A	No
Liu et al. [5] (GNN)	No	N/A	No
Yang et al. [6] (RL dispatch)	No	No	Limited
Proposed framework	Yes	Yes	Yes (IEEE 33-bus)

III. PROPOSED DIGITAL TWIN ARCHITECTURE

The proposed framework has five connected layers: (1) the physical grid layer - PV/wind assets, storage, and load nodes with IoT sensors; (2) a data-acquisition layer streaming telemetry via SCADA/IoT gateways; (3) the digital twin core, holding a physics-informed, continuously updated network state model; (4) an AI optimisation engine; and (5) a scenario-simulation sandbox that tests candidate control actions before dispatch, closing the feedback loop.

A. AI Optimisation Engine

The AI engine combines three cooperating models, summarised in Table II. An LSTM/Transformer network forecasts short-horizon PV and wind output. A graph neural network encodes feeder topology for robust state estimation and fault localisation. A DRL agent, trained in the simulation sandbox, computes dispatch actions - storage charge/discharge, curtailment set-points, tap-changer positions - to maximise renewable utilisation within voltage and thermal limits.



TABLE II COMPARISON OF AI TECHNIQUES EMBEDDED WITHIN THE DIGITAL TWIN

Technique	Primary Function	Accuracy / Gain	Comp. Cost
LSTM / Transformer	Short-term solar & wind forecasting	92-97% (1h ahead)	Medium
Graph Neural Network	Topology-aware state estimation & fault localisation	95% node-state accuracy	High
Deep RL	Real-time dispatch / storage optimisation	15-30% cost reduction	High (offline)
Autoencoder / GAN	Synthetic scenario generation & anomaly detection	90%+ anomaly recall	Medium

B. Scenario Simulation Sandbox

Before any control action reaches the physical grid, it first runs inside a high-fidelity replica within the digital twin. This sandbox checks the predicted effect on voltage, loading, and frequency, rejecting or adjusting actions that would breach operating limits. Actions that pass go to the grid actuators; those that fail return to the DRL agent with a penalty signal.

C. Data Synchronisation and Latency Management

The digital twin core stays synchronised with the physical grid via a publish-subscribe telemetry bus, keeping measurement staleness under two seconds for voltage/current phasors and under fifteen seconds for weather telemetry. A state-reconciliation routine periodically corrects drift using a Kalman-filter correction step.

D. Security Considerations

Since the digital twin both receives operational data and issues control recommendations, it is treated as a critical cyber-physical asset. The architecture uses mutually authenticated telemetry channels, role-based access control for sandbox overrides, and anomaly detection on the AI engine's own inputs and outputs.

E. Mathematical Formulation of the Dispatch Problem

The DRL agent's dispatch problem is a Markov decision process: state s_t covers forecasted generation, storage state-of-charge, and GNN-estimated bus voltages; action a_t covers storage charge/discharge power and curtailment set-points; reward $r_t = -(w_1 * C_{cost}(t) + w_2 * C_{curtail}(t) + w_3 * C_{voltage}(t))$. Policy $\pi_{\theta}(a_t | s_t)$ is optimised to maximise expected discounted return via proximal policy optimisation, using discount factor $\gamma = 0.98$.

The GNN state estimator solves a semi-supervised node-regression problem over feeder graph $G = (V, E)$, with nodes as buses and edges as line segments, aggregating neighbour messages across two graph-convolution layers. The forecasting model minimises mean-squared-error loss, with attention weighting recent time-steps and similar-weather historical days more heavily.

IV. METHODOLOGY

A modified IEEE 33-bus radial feeder was augmented with 6 MW aggregate PV, 3 MW wind, and two 2 MWh battery storage units to emulate a high-renewable network, as in Table III. One year of synthetic irradiance, wind-speed, and load data at 5-minute resolution was generated from historical weather profiles with added stochastic noise.

TABLE III TEST-SYSTEM AND TRAINING CONFIGURATION

Parameter	Value	Notes
PV capacity	6 MW (18 nodes)	Rooftop + feeder arrays
Wind capacity	3 MW (4 nodes)	Small onshore turbines
Battery storage	2 x 2 MWh	Li-ion, 90% RTE
Feeder topology	IEEE 33-bus radial	12.66 kV base
Data resolution	5 minutes	1yr synthetic + weather



DRL training	5,000 episodes	PPO
Forecast horizon	1-6 hours	Re-forecast every 15 min

The forecasting sub-model was trained with a sliding-window supervised approach on a 70/15/15 split. The GNN state estimator trained on partially observed data, masking 30% of nodes at random. The DRL dispatch agent trained via PPO in the sandbox for 5,000 episodes, with rewards penalising voltage deviation, curtailment, and cost.

All models were benchmarked against a rule-based conventional EMS on the same feeder and data set. Six metrics were evaluated - renewable utilisation, curtailment reduction, voltage stability index, fault response time, forecast RMSE, and operational cost.

A. Evaluation Metrics

Renewable utilisation is the share of available PV/wind energy delivered to load without curtailment. Voltage stability index is the share of time all bus voltages stay within the $\pm 5\%$ statutory band. Fault response time runs from disturbance onset to correct isolation of the affected section. Forecast RMSE compares predicted and actual generation, and operational cost combines energy-purchase, curtailment-penalty, and storage-degradation costs into one \$/MWh figure.

B. Computational Environment

Training ran on a workstation with a 16-core CPU and one 24 GB GPU. Table IV reports training and inference times per AI component, confirming that run-time inference for all three stays well within the sub-second budget needed for real-time dispatch.

TABLE IV COMPUTATIONAL PERFORMANCE OF THE AI COMPONENTS

Component	Offline Training	Inference Latency	Update Freq.
LSTM / Transformer	~3.5 h	<50 ms	Daily
GNN state estimator	~6.0 h	<120 ms	Weekly
DRL dispatch agent	~14.0 h	<80 ms	Monthly
Simulation sandbox check	N/A	<200 ms	Every cycle

V. RESULTS AND DISCUSSION

Table V lists the full quantitative results from the case study. The framework raised renewable utilisation from 72.0% to 91.0% and cut curtailment by 28 points. Fault response time dropped from 850 ms to 270 ms, driven by the GNN-based estimator localising affected sections faster, while forecast RMSE fell by 59.4%.

TABLE V SIMULATION RESULTS: CONVENTIONAL EMS VS. PROPOSED DT-AI FRAMEWORK

Performance Metric	Conventional EMS	Proposed DT-AI	Improvement
Renewable utilisation (%)	72.0	91.0	+19.0 pp
Curtailment reduction (%)	18.0	46.0	+28.0 pp
Voltage stability index (%)	74.0	93.0	+19.0 pp
Fault response time (ms)	850	270	-68.2%
Forecast RMSE (kW)	42.6	17.3	-59.4%
Operational cost (\$/MWh)	58.2	44.7	-23.2%

The 23.2% operational cost cut mainly comes from reduced curtailment and more efficient storage cycling from the DRL agent, which learns to charge during surplus generation and discharge during high marginal cost. Voltage stability index rose from 74.0% to 93.0%, showing the sandbox effectively filters out actions that would push bus voltages toward their limits.

These gains come with added offline training and compute cost for the GNN and DRL components, suggesting an edge-cloud hybrid deployment - lightweight inference at the substation edge, periodic retraining in the cloud - suits large feeders with limited local compute.



A sensitivity analysis found performance degrades gracefully under greater measurement sparsity: renewable-utilisation gains fall from 19 to about 12 points when 50% of feeder nodes lack direct instrumentation, underscoring the value of the GNN's topology-aware imputation.

VI. IMPLEMENTATION CHALLENGES AND FUTURE WORK

Several practical hurdles remain before wide-scale deployment is feasible. Data quality and completeness vary widely across utilities, and legacy SCADA infrastructure may not support the measurement resolution assumed here. The compute and energy cost of continually retraining GNN and DRL components must be weighed against operational gains, especially for smaller utilities. Regulatory and certification frameworks for AI-driven autonomous control in critical infrastructure remain immature, and utilities may initially need a human-in-the-loop approval step.

A limitation of this study is its reliance on a single feeder topology and synthetic weather data; field validation across multiple feeders with real telemetry is needed to confirm generalisability. The reward-weighting coefficients were tuned on a validation split of the same synthetic data set - performance under materially different tariffs or weather has not been verified. Future work will extend the framework to multi-feeder and transmission-level twins, add federated learning across utilities, and test resilience against adversarial telemetry manipulation.

VII. CONCLUSION

This paper introduced an AI-based digital twin framework for renewable-rich distribution grids, unifying LSTM/Transformer forecasting, GNN-based state estimation, and DRL-driven dispatch within a continuously synchronised virtual replica and safety-validation sandbox. Unlike prior work that treats these components in isolation, the proposed architecture couples them inside a single closed-loop system so that every control decision is grounded in live forecasts and state estimates before release to the physical grid.

Case-study results on a modified IEEE 33-bus feeder demonstrate the practical value of this integration: renewable utilisation rose from 72.0% to 91.0%, curtailment fell markedly, fault response time dropped by 68.2%, and operational cost declined by 23.2% relative to a conventional EMS. The sensitivity analysis shows the framework degrades gracefully under sparse instrumentation, and the sandbox reliably filters out unsafe actions before they reach the field.

These findings position AI-enhanced digital twins as a viable path toward self-optimising, resilient distribution grids. Future work will extend the framework to multi-feeder and transmission-level twins, incorporate federated learning, and evaluate robustness against adversarial telemetry manipulation using field-deployed sensor data.

REFERENCES

- [1] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary Perspectives on Complex Systems*, Springer, 2024, pp. 85-113.
- [2] F. Tao et al., "Digital twin-driven product design, manufacturing and service with big data," *Int. J. Adv. Manuf. Technol.*, vol. 131, pp. 3563-3576, 2024.
- [3] H. Wang et al., "A review of digital twin applications in smart grids," *IET Renew. Power Gener.*, vol. 18, no. 3, pp. 412-430, 2024.
- [4] Y. Wu, M. Wu, and X. Chen, "Short-term renewable generation forecasting using Transformer networks," *IEEE Trans. Sustain. Energy*, vol. 15, no. 1, pp. 210-222, 2024.
- [5] Z. Liu et al., "Graph neural networks for power system state estimation," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 1889-1901, 2024.
- [6] T. Yang, L. Zhao, and W. Li, "Reinforcement learning for microgrid energy management: A 2024 review," *Appl. Energy*, vol. 356, 122350, 2024.
- [7] S. Singh and A. Kumar, "Machine learning based fault detection in distribution networks with renewable integration," *IET Gener. Transm. Distrib.*, vol. 18, no. 6, pp. 987-1001, 2024.
- [8] IEC, "Digital twins for industrial and energy applications," IEC White Paper, 2024.
- [9] R. Patel and D. Shah, "Deep learning approaches for smart grid optimisation: A survey," *Renew. Sustain. Energy Rev.*, vol. 191, 114120, 2024.



- [10] C. Zhang and P. Rajagopal, "Federated learning for distributed energy resource forecasting," IEEE Trans. Smart Grid, vol. 15, no. 3, pp. 2541-2553, 2024.
- [11] N. Kumar et al., "Cybersecurity of digital twins in critical infrastructure: A survey," Comput. Secur., vol. 138, 103654, 2024.
- [12] J. Lopez and E. Martins, "Kalman-filter based state reconciliation for real-time digital twins," Electric Power Syst. Res., vol. 229, 110123, 2024

