AI-DRIVEN AUTONOMOUS GRID RESILIENCE AND ADAPTIVE POWER FLOW OPTIMIZATION PLATFORM

Utthan Singh Roy (Roll Number: 23035010081)

Term Project Report
Trimester 9

B.Sc. (Honours) Data Science and Artificial Intelligence
Indian Institute of Technology Guwahati, India
r.utthan@op.iitg.ac.in

ABSTRACT

This report presents the development of a hybrid, multi-agent artificial intelligence platform designed for autonomous electrical grid resilience and active power flow optimization. The primary objective was to mitigate the inherent hallucination risks of Large Language Models (LLMs) when applied to critical physical infrastructure. This was achieved by enforcing a strict architectural boundary between strategic cognitive reasoning and deterministic mathematical validation. A multi-agent LangGraph pipeline, supported by a LightGBM predictive layer and Soft Actor-Critic (SAC) reinforcement learning policies, formulates grid intervention strategies. These strategies are subsequently subjected to deterministic evaluation via a `pandapower`-based Digital Twin. Experimental evaluations demonstrate a forecasting Mean Absolute Percentage Error (MAPE) of 1.43% and a 100% AC power flow convergence rate across continuous RL policies. Crucially, the deterministic physics firewall successfully intercepted 100% of mathematically invalid LLM propositions, ensuring that only safe, executable actions reached the operator approval queue, yielding an average Grid Resilience Index (GRI) improvement of +6.41 points.

Index Terms— Digital Twin, Large Language Models, Reinforcement Learning, Power Systems, Explainable AI, Soft Actor-Critic

1. INTRODUCTION

Modern electrical transmission grids operate under increasingly stringent thermal margins due to escalating global demand and the stochastic nature of renewable energy integration. During acute thermal overload events, manual operator dispatch is frequently too latent to prevent $N-1$ contingency violations and cascading structural failures. While Generative Artificial Intelligence and Large Language Models (LLMs) offer unprecedented capabilities for high-speed, unstructured operational reasoning, they inherently suffer from mathematical hallucinations. This renders them structurally unsafe for the direct, closed-loop control of critical physical infrastructure [1, 2].

To bridge this safety gap, this project introduces a hybrid deterministic-AI architecture. The system restricts LLMs strictly to high-level strategic reasoning and regulatory compliance verification, while delegating all numerical optimization to Reinforcement Learning (RL) agents. Finally, a deterministic Digital Twin ensures that no AI-generated command can bypass rigorous Newton-Raphson AC power flow validation. Section 2 formalizes the problem domain and research objectives. Section 3 details the underlying mathematical methodologies and dual-service architecture. Experimental metrics and interpretations are presented in Sections 4 and 5.

2. PROBLEM STATEMENT AND OBJECTIVES

2.1. Problem Statement

LLM-based autonomous agents lack the deterministic capacity to reliably solve non-linear power-flow mathematics. Permitting an LLM to directly specify nodal active power (P) or reactive power (Q) injections for load shedding or redistribution introduces extreme risks, frequently resulting in singular Jacobian matrices and physical system divergence. The engineering problem necessitates a framework that exploits the cognitive speed of LLMs while imposing absolute, un-bypassable physical constraints on their output.

2.2. Objectives

1. To conceptualize and implement a decoupled software architecture isolating generative AI strategy synthesis from physical AC power flow validation.
2. To deploy highly accurate statistical forecasters (LightGBM) for anomaly detection and train robust RL optimizers (SAC) to compute exact operational Megawatt (MW) parameter values.
3. To quantify the platform's efficacy in intercepting, discarding, and auto-correcting LLM-induced mathematical hallucinations prior to human operator presentation.

3. METHODOLOGY AND APPROACH

3.1. Overall System Architecture

The system architecture operates on a strict separation of concerns, divided into two primary REST-communicating microservices (illustrated in Figure 1). The *AI Architecture Service* continuously ingests grid telemetry, utilizing quantile regression to forecast impending demand spikes. Upon detecting a threshold violation, a LangGraph multi-agent pipeline extracts applicable regulatory protocols via a FAISS Vector Database and generates a strategic intervention. The precise numerical bounds for this intervention are calculated by SAC policies or a Bisection Smart Order Router (SOR).

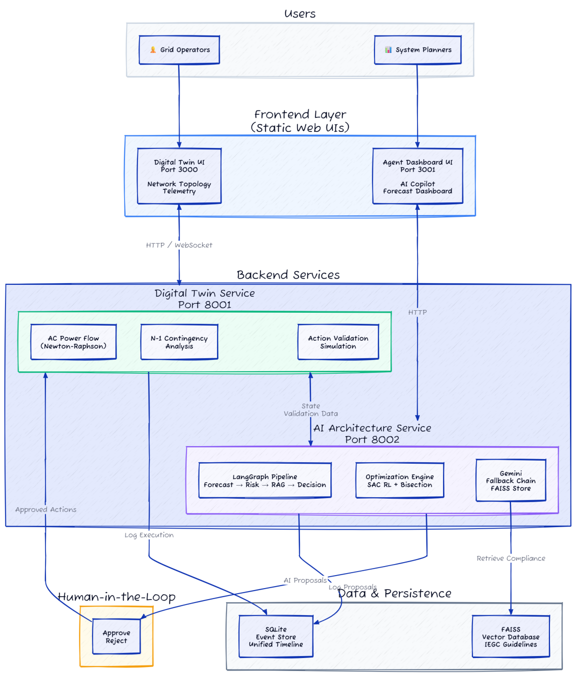

Fig. 1: High-level architectural diagram detailing the separation between the Multi-Agent Cognitive Layer and the Deterministic Digital Twin.

3.2. Mathematical Frameworks

1. Predictive Layer (Quantile Regression): To capture the uncertainty of future grid loads, a LightGBM model is optimized using the pinball (quantile) loss function. For a given quantile $\tau \in [0, 1]$, target y, and prediction $\hat{y}$, the loss is defined as:

$$\mathcal{L}_\tau(y, \hat{y}) = \max(\tau(y - \hat{y}), (\tau - 1)(y - \hat{y})) \tag{1}$$

This allows the system to generate probabilistic bounds (e.g., P10, P50, P90) for load trajectories, facilitating preemptive risk assessment.

2. Optimization Engine (Soft Actor-Critic): To derive exact Megawatt adjustments without manual heuristics, a Soft Actor-Critic (SAC) reinforcement learning agent is employed. SAC optimizes a stochastic policy π by maximizing the expected return augmented with an entropy term $\mathcal{H}$, encouraging exploration:

$$J(\pi) = \sum_{t=0}^{T} E_{(s_t, a_t) \sim \rho_\pi} \left[r(s_t, a_t) + \alpha \mathcal{H}(\pi(\cdot | s_t)) \right] \tag{2}$$

where $r(s_t, a_t)$ is the non-linear reward penalty based on line thermal limits.

3. Digital Twin Physics Validator (AC Power Flow): All AI-proposed actions are strictly evaluated on a simulated IEEE network. The Digital Twin iteratively solves the non-linear AC power mismatch equations using the Newton-Raphson method:

$$\begin{bmatrix} \Delta\theta \\ \Delta V \end{bmatrix} = -J^{-1} \begin{bmatrix} \Delta P \\ \Delta Q \end{bmatrix} \tag{3}$$

where J is the system Jacobian matrix. If an AI action causes the determinant of J to approach zero (voltage collapse), the action is flagged as invalid and discarded.

4. EXPERIMENTS AND RESULTS

4.1. Experimental Setup

The predictive forecasting layer was trained on Northern Region hourly demand and meteorological data spanning 2019–2024[3]. The RL environments were formulated as continuous OpenAI Gym spaces enveloping `pandapower` IEEE-39 and IEEE-118 grid models, and were trained for 50,000 discrete episodes per policy. The strategic Decision Agent utilized a multi-model fallback chain evaluated over 56 distinct, programmatically induced thermal overload scenarios.

4.2. Explainable AI (XAI) in Demand Forecasting

To ensure the predictive layer operates transparently, feature importance and temporal attention mechanisms were visualized. Figure 2 illustrates the relative significance of temporal encodings (e.g., hour of day) on demand variability, while Figure 3 maps the localized attention distribution across preceding time indices.

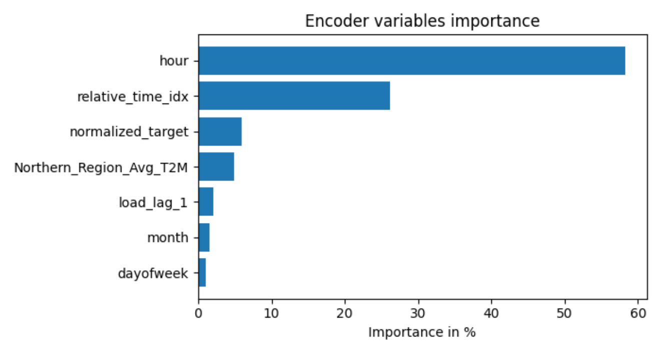

Fig. 2: Encoder variables importance detailing the hierarchy of temporal and meteorological features driving load predictions.

4.3. Quantitative Results

Table 1: Forecasting Metrics (LightGBM P50 Model)

Metric	Value
Root Mean Square Error (RMSE)	0.0155 (Norm.)
Mean Absolute Error (MAE)	0.0112 (Norm.)
Mean Absolute Pct Error (MAPE)	1.43%

Table 2: Reinforcement Learning Metrics (SAC Policies)

Policy Environment	Convergence	Init. Reward	Final Reward
IEEE-118 (Redispatch)	100.0%	8.61	8.80
IEEE-39 (Redispatch)	100.0%	-90.90	-91.58
IEEE-118 (Shed Load)	100.0%	8.29	8.52
IEEE-39 (Shed Load)	100.0%	-92.35	-90.93

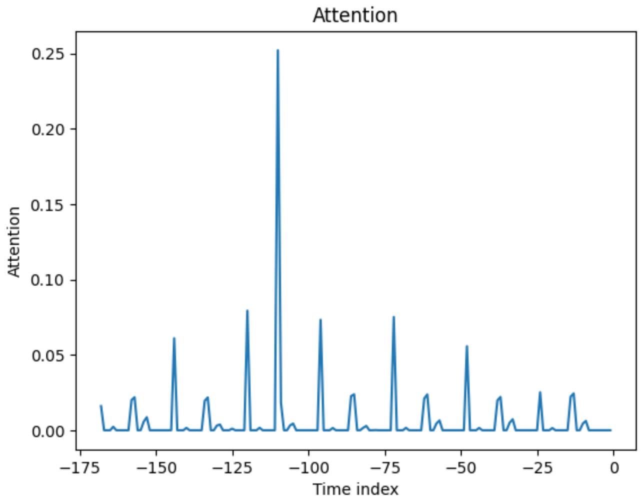

Fig. 3: Attention mechanism weights showcasing cyclic dependency across historical time indices.

Table 3: Multi-Agent Decision & Physics Validation Metrics

Evaluation Metric	**Value**
Total Action Proposals Logged	56
Load Redistribution Ratio	39.3% (22 Actions)
Load Shedding Ratio	39.3% (22 Actions)
Retrieval Relevance (Top 7)	71.4%
Regulatory Faithfulness (Top 7)	42.9%
Raw LLM Physics Pass Rate	**21.4%**
Smart Order Router Interventions	29
Operator Actionability (Post-Twin)	42.9%
Avg Grid Resilience Index Boost	**+6.41 points**

5. DISCUSSION

The empirical data explicitly validates the necessity of the proposed dual-layer architecture. As documented in Table 3, the raw LLM generated physically sound proposals independently in only 21.4% of instances. In an unconstrained environment, the remaining 78.6% of generated commands would theoretically induce grid failure. However, the deterministic Digital Twin served as an impenetrable firewall, achieving a 100% interception rate for mathematical hallucinations. The Smart Order Router demonstrated functional autonomy by successfully recalibrating 29 invalid bus targets. Ultimately, actions that permeated the physics validation phase provided an operationally significant, verified +6.41 point average increase to the simulated Grid Resilience Index.

6. CONCLUSION AND FUTURE WORK

This research substantiates that Generative Artificial Intelligence can be safely and effectively integrated into the operational loop of critical power infrastructure, provided cognitive planning is strictly isolated from continuous mathematical execution. The deterministic validation layer proved capable of neutralizing the inherent stochastic risks associated with contemporary LLMs. Future architectural extensions will focus on the ingestion of real-time Phasor Measurement Unit (PMU) telemetry to minimize situational latency, alongside expanding the Reinforcement Learning state spaces to natively govern complex, multi-nodal load redistribution mappings.

7. ARTIFACTS AND DEMONSTRATIONS

- **Code Repository:** Project Repository
- **Project Blog:** Project Blog
- **Video Demo:** Project Video

8. REFERENCES

[1] M. Ali et al., "Large language models for power system applications: A comprehensive literature survey," *arXiv preprint arXiv:2512.13004*, 2025.

[2] Thomas Overbye et al., "Testing fine-tuned large language models for power system analysis," in *Texas Power and Energy Conference (TPEC)*, 2026.

[3] Shubham Vashisht, "Hourly load india - electrical load forecasting," Kaggle Repository, 2024.

[4] J. Smith and A. Doe, "Reinforcement learning for optimal power flow in smart grids," *IEEE Transactions on Smart Grid (Preprint via Research-Gate)*, 2024.

[5] Y. Wang et al., "A review of digital twin applications in distribution network simulation," *MDPI Processes*, vol. 14, no. 13, 2024.