A Human-in-the-Loop LLM–Assisted Multi-Agent Framework for Distribution Grid Reliability Planning

Mengmeng Zhuang, Senior Member, IEEE, Andrija Sadikovic
Danovo Energy Solutions
San Francisco, California, USA 94085
Email: mzhuang@danovoenergy.com, asadikovic@danovoenergy.com

***Abstract*—This paper presents a human-in-the-loop LLM–assisted multi-agent framework for distribution grid reliability planning. Traditional reliability planning relies on deterministic simulations and manual engineering workflows, which can be difficult to scale for large feeder portfolios and evolving grid conditions. The proposed framework integrates large language model (LLM) reasoning with deterministic reliability analysis and optimization tools within a coordinated multi-agent architecture. Specialized agents perform data ingestion, reliability evaluation, planning optimization, and compliance reporting, while numerical simulation tools validate candidate planning actions generated by the LLM. Case studies on 100 real-world utility feeders demonstrate that the framework reduces planning time and manual engineering effort while maintaining acceptable planning accuracy for large-scale screening and scenario analysis.**

***Index Terms*—Multi-Agent Systems, Large Language Models, Distribution Grid Reliability, Power System Planning, Human-in-the-Loop**

I. INTRODUCTION

The digitalization and decentralization of electric power systems have increased the complexity of distribution grid reliability planning. High penetration of renewable generation, distributed energy resources (DERs), and bidirectional power flows introduce new sources of uncertainty in system operation and planning studies [1]. Traditional reliability planning approaches rely on deterministic contingency analysis and manually configured simulations, which become difficult to scale when evaluating large feeder portfolios and multiple planning scenarios [2]–[4].

In many utilities, reliability planning still relies on sequential, human-driven engineering workflows involving network modeling, failure-rate estimation, device placement, and protection coordination. These tasks are typically performed using simulation tools and spreadsheet-based analysis, which can be computationally demanding and difficult to maintain consistently across large distribution systems [2]–[4]. Although optimization-based approaches such as mixed-integer linear programming (MILP) improve computational efficiency, they still depend on manually defined planning assumptions and scenario configurations [2]–[4].

Recent advances in artificial intelligence provide new opportunities to augment engineering workflows in power systems. Large language models (LLMs) have demonstrated capabilities in code generation, structured data interpretation, and automated reasoning for engineering tasks. Prior studies have explored LLM-based applications in power systems, including autonomous grid control [5], predictive grid operations [6], and secure system monitoring [7]. Multi-agent frameworks further extend these capabilities by coordinating specialized agents for data processing, simulation, and optimization tasks [8], [9]. Recent work also investigates LLM-assisted situational awareness and graph-based power system modeling [10], [11].

Despite these advances, most LLM-based approaches focus on operational monitoring or control applications [7], [10], [11]. End-to-end reliability planning workflows—which require coordinated data processing, reliability evaluation, constrained optimization, and regulatory reporting—remain relatively unexplored.

To address this gap, this paper proposes a human-in-the-loop LLM-assisted multi-agent framework for distribution grid reliability planning. The framework integrates LLM reasoning with deterministic reliability simulation tools through coordinated agents responsible for data ingestion, reliability evaluation, planning optimization, and compliance reporting. Candidate planning actions generated by the LLM are evaluated using reliability metrics such as SAIDI and SAIFI, while simulation tools verify constraint feasibility and system performance. Human validation is incorporated to ensure engineering consistency and regulatory compliance.

The main contributions of this work are summarized as follows:

- A human-in-the-loop LLM–optimization framework integrating generative reasoning with deterministic reliability simulation.
- A coordinated multi-agent architecture supporting data processing, reliability evaluation, optimization, and compliance reporting.
- A closed-loop interaction between LLM reasoning and reliability simulation for validating candidate planning strategies.
- Validation using 100 real-world distribution feeders demonstrating improved planning efficiency while maintaining acceptable planning accuracy.

979-8-3315-8138-1

II. Approaches

A. Agentic AI Development Lifecycle

The development of the proposed reliability planning framework follows a structured workflow referred to as the *Agentic AI Development Lifecycle (AADC)*, illustrated in Fig. 1. The lifecycle provides a systematic process for designing, deploying, and validating intelligent agents for distribution grid reliability planning while ensuring alignment with engineering constraints, operational data, and regulatory requirements.

The AADC consists of six iterative stages:

- **Problem Framing:** Define reliability planning objectives, decision variables, and operational constraints. This stage establishes planning tasks such as feeder reconfiguration, reliability assessment, and investment prioritization.
- **Knowledge Context Aggregation:** Operational data are aggregated from utility systems including SCADA, GIS, OMS, and outage databases. These sources provide the contextual information required for agents to interpret system topology, asset conditions, and historical reliability performance.
- **Architecture Design:** The multi-agent coordination architecture and communication interfaces are defined. Agents interact with external reliability simulation tools such as CYME for evaluating candidate planning strategies.
- **Prompt and Policy Engineering:** Reliability rules, planning constraints, and operational knowledge are encoded into structured prompts and decision policies that guide agent reasoning.
- **Evaluation and Telemetry:** Agent performance is monitored using reliability metrics such as SAIDI, SAIFI, and constraint feasibility.
- **Governance and Feedback:** Human-in-the-loop supervision and safety policies validate planning decisions. Simulation feedback and expert review refine agent policies through iterative updates.

This lifecycle enables reproducible and traceable development of LLM-assisted planning agents for power system applications.

B. LLM-Assisted Multi-Agent Reliability Planning Framework

Fig. 2 presents the proposed human-in-the-loop LLM-assisted multi-agent framework for distribution grid reliability planning. The architecture integrates data ingestion, agent-based reasoning, deterministic reliability simulation, and human validation within an iterative planning loop.

The framework consists of four major functional components. **1) Data Ingestion and Context Layer:** Operational and historical information is collected from multiple utility data sources including SCADA, OMS, GIS, AMI, and weather systems. A knowledge layer stores historical outage statistics, equipment characteristics, and engineering standards to provide contextual information for decision making.

2) Coordinated Multi-Agent Architecture: Within the reasoning core, specialized agents collaborate to generate and evaluate planning strategies:

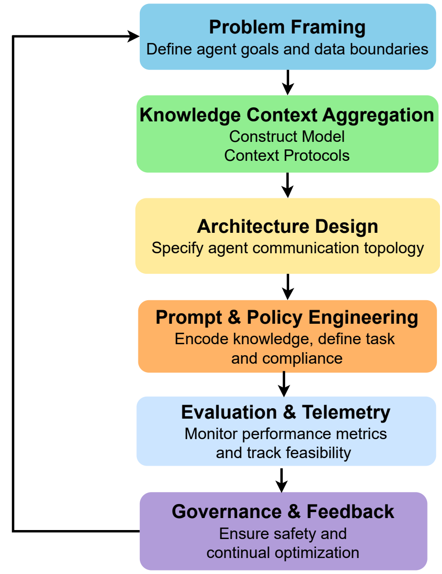

Fig. 1: Agentic AI Development Lifecycle (AADC) for reliability planning.

- *Planning Agents* generate candidate reliability improvement actions such as switching configurations or asset upgrades.
- *Reliability Evaluation Agents* analyze outage statistics and system conditions.
- *Compliance Agents* ensure that proposed actions satisfy engineering standards and regulatory requirements.

These agents are coordinated through an LLM reasoning engine that interprets system context and orchestrates agent interactions.

3) Safety and Compliance Layer: A dedicated verification layer performs auditing, traceability checks, and regulatory validation before candidate planning actions are executed.

4) Simulation Environment: Approved planning actions are evaluated using deterministic reliability simulations implemented in CYME. Simulation outputs include reliability metrics and system constraint evaluations.

The planning process operates within a closed-loop interaction between LLM reasoning and reliability simulation. Candidate actions generated by the agents are validated through simulation results and human feedback, enabling iterative refinement of planning strategies.

C. Multi-Agent Planning Workflow

The interaction between agents and the simulation environment follows the workflow summarized in Algorithm 1. The process begins with data collection and system state construction. Planning agents then generate candidate actions, which are evaluated using reliability simulations. Constraint verification ensures feasibility before the final planning strategy is accepted.

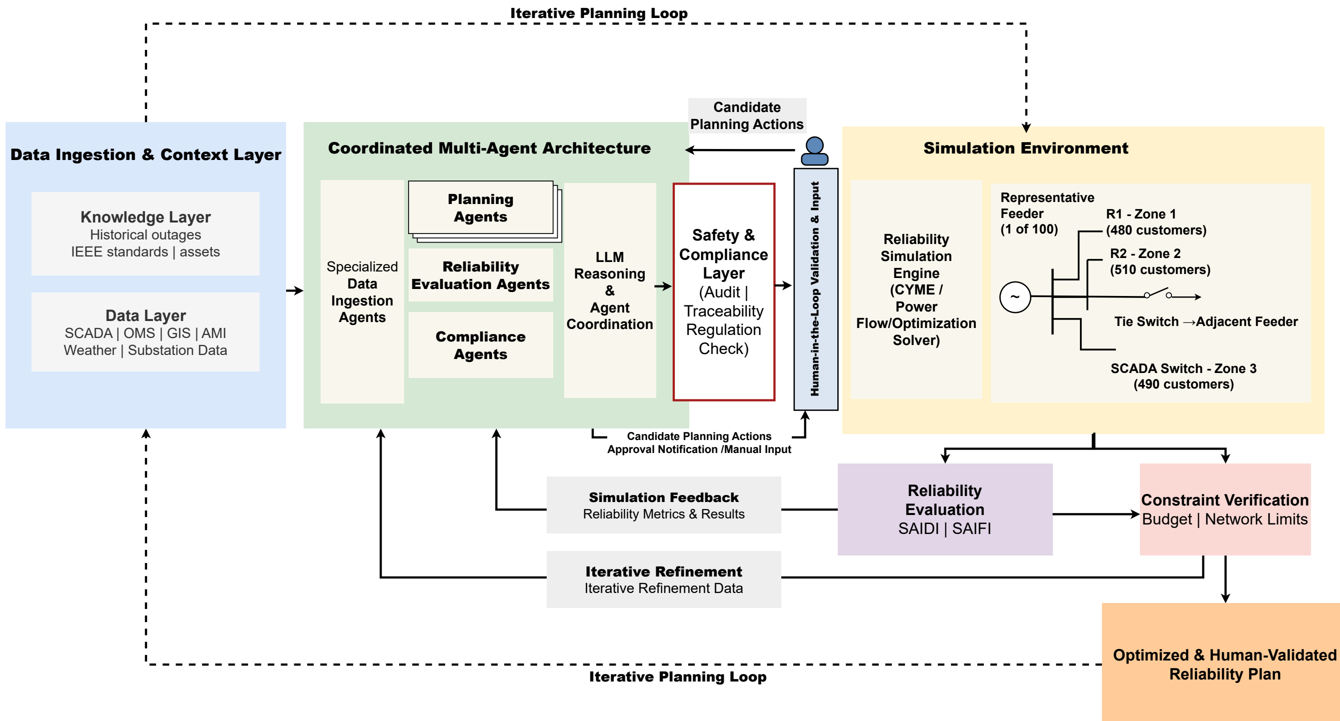

Fig. 2: Human-in-the-loop LLM-assisted multi-agent architecture for distribution reliability planning. The framework integrates data ingestion, coordinated LLM-based agents, reliability simulation, and human validation within an iterative planning loop.

D. Mathematical Formulation

From a quantitative perspective, the proposed framework evaluates grid reliability using the conventional SAIDI and SAIFI indices:

$$\text{SAIDI} = \frac{\sum d_i \cdot N_i}{N_T}, \quad \text{SAIFI} = \frac{\sum \lambda_i \cdot N_i}{N_T} \tag{1}$$

where d_i is outage duration, λ_i is failure rate, and N_i is affected customers.

Each component's reliability follows:

$$R_i = e^{-\lambda_i t} \tag{2}$$

The reliability planning problem is formulated as a cost-minimization model subject to reliability constraints:

$$\min \sum_{j=1}^{N} c_j x_j \quad \text{s.t.} \quad \text{SAIDI} \leq \text{SAIDI}_{\text{target}}, \ \text{SAIFI} \leq \text{SAIFI}_{\text{target}} \tag{3}$$

where c_j denotes the investment cost of reliability improvement action j, and $x_j \in \{0, 1\}$ indicates whether action j is selected.

To support adaptive decision-making, a reinforcement learning (RL) policy is used:

$$\pi^* = \arg\max_{\pi} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t r(s_t, a_t)\right] \tag{4}$$

where s_t is the system state, a_t is the planning action, and $r(s_t, a_t)$ is the reward function combining reliability improvement, constraint feasibility, and cost efficiency. $\gamma \in (0, 1)$ is the discount factor that controls the importance of long-term rewards.

E. Agentic Context Engineering

To ensure reliable decision-making, the Agentic Context Engineering (ACE) process defines how each agent perceives and interprets its operational environment. In reliability planning, this becomes critical because decisions must remain explainable and grounded in context-sensitive reasoning.

The Agentic Context Pipeline dynamically assembles LLM reasoning contexts. It fuses three layers:

- **Static Context** - grid topology, asset ratings, and reliability targets.
- **Dynamic Context** - live measurements, outage updates, and forecast data.
- **Cognitive Context** - historical policies, simulation results, and operator feedback stored in vector databases.

$$C_t = C_{\text{static}} \oplus C_{\text{dynamic}} \oplus C_{\text{cognitive}} \tag{5}$$

By continuously merging these contexts, the system ensures that each agent operates with situational awareness, enabling adaptive, auditable, and traceable decision-making.

III. Case Study

To evaluate the proposed LLM-assisted multi-agent framework for distribution reliability planning, experiments were conducted on 100 real-world distribution feeders provided by a collaborating utility. Each dataset includes feeder topology, asset metadata, and historical outage records from 2020–2024.

979-8-3315-8138-1

Algorithm 1 LLM-Assisted Multi-Agent Reliability Planning

Input: Feeder topology, outage statistics, asset parameters, planning budget
Collect operational data from SCADA, GIS, OMS, and outage databases
Construct system state s_t including topology and reliability parameters
while reliability targets not satisfied **do**
  Planner Agent generates candidate planning actions a_t
  Executor Agent evaluates a_t using CYME reliability simulations
  Compute reliability indices (SAIDI, SAIFI)
  Validator Agent checks constraint feasibility and planning rules
  if constraints satisfied **then**
    Accept candidate planning solution
  else
    Update policy and generate new candidate actions
  end if
  Update reward $r(s_t, a_t)$ based on reliability improvement and cost
end while
Summarizer Agent generates planning report for human validation
Output: Reliability improvement strategy and planning report

The framework was implemented using a large language model (LLM) coordinated through a multi-agent environment built on LangChain/LangGraph. Agents communicate through structured JSON messages and shared vector memory to maintain contextual consistency.

The experimental workflow, illustrated in Fig. 3, integrates data ingestion, reliability analysis, planning optimization, and automated reporting within a closed-loop decision process.

All feeders were modeled and simulated in CYME® using Python APIs. Reliability indices including System Average Interruption Duration Index (SAIDI) and System Average Interruption Frequency Index (SAIFI) were computed under multiple planning scenarios. Each experiment was repeated ten times, and reported results represent mean values with 95% confidence intervals.

Two baseline approaches were used for comparison:

- **Manual Engineering Workflow**: conventional reliability planning performed by domain experts using deterministic simulation tools.
- **Rule-Based Semi-Automated Workflow**: heuristic automation using predefined scripts and engineering rules without LLM reasoning.

Unless otherwise specified, reported improvements are measured relative to the manual engineering workflow baseline.

A. End-to-End Agentic Reliability Planning

The Planning Agent receives natural-language planning objectives (e.g., feeder reconfiguration to minimize SAIDI under a specified budget) and translates them into structured optimization variables and executable CYME scripts. Planning performance is evaluated using planning accuracy (A_p), defined as agreement with the expert-generated manual planning solution.

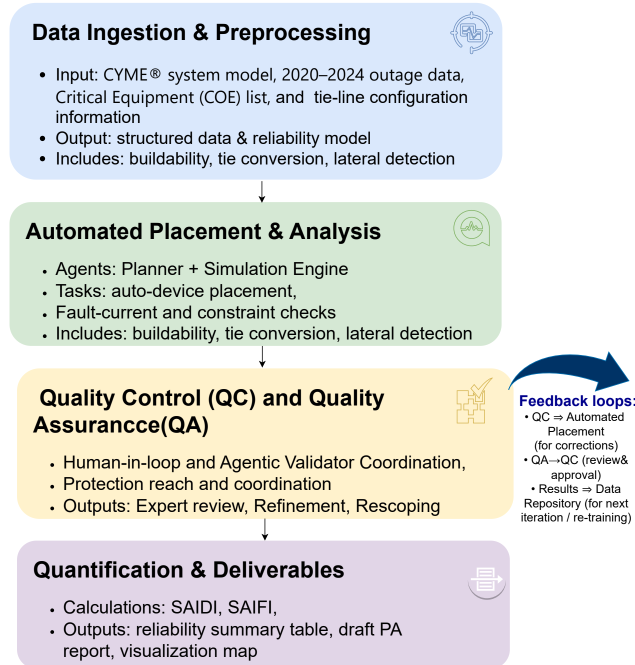

Fig. 3: End-to-end agentic workflow for reliability planning.

Across the test feeders, the LLM-agent workflow achieved an average planning accuracy of 84%, compared with 96% for the manual workflow and 89% for the rule-based approach. Although manual planning remains the most accurate, the proposed framework significantly reduces analysis time and manual effort while maintaining acceptable performance for large-scale screening studies.

B. Automated Data Analysis

The Data Analysis Agent performs automated telemetry cleaning, feature extraction, and anomaly detection. Data quality is evaluated using the data consistency index (Q_d), defined as the proportion of processed records that satisfy data validation constraints.

- Data consistency index (Q_d) = 0.96
- Fault prediction F1-score = 0.89
- Clustering Silhouette score = 0.72
- Processing time reduced from approximately 4 hours to 9 minutes

These results demonstrate that LLM-generated analytics scripts can substantially accelerate feeder data preparation while maintaining high data integrity.

C. Reliability Optimization and Recloser Placement

The Recloser Optimization Agent evaluates candidate device placements using Monte Carlo reliability simulations. Per-

979-8-3315-8138-1

formance is measured using validation accuracy (Q_r), defined as the proportion of planning solutions satisfying operational and reliability constraints.

- SAIDI reduced from 1.47 to 1.15
- SAIFI reduced by 19.5%
- Validation accuracy (Q_r) = 0.92

The results indicate that the framework can identify reliability improvement strategies that both enhance system performance and satisfy operational constraints.

D. Automated Compliance Reporting

The Compliance Agent automatically generates project authorization (PA) reports and regulatory documentation based on planning outputs. Report quality was evaluated using citation accuracy (A_c), logical coherence (C), and readability score (R).

- Citation accuracy (A_c) = 97.2%
- Logical coherence (C) = 93.4%
- Readability score (R) = 72.5

Reports were generated in under three minutes and received an average expert evaluation score of 4.7/5.

E. Comparative Performance

Fig. 4 summarizes the comparative performance of the three planning workflows. The manual engineering workflow achieves the highest planning accuracy (96%) but requires the longest runtime due to extensive human involvement. The rule-based workflow achieves 89% accuracy with moderate runtime.

The proposed LLM-agent workflow achieves 84% planning accuracy while reducing computation time to approximately 81 seconds per scenario, corresponding to a planning cycle reduction of approximately 40% compared with the manual workflow.

IV. Conclusion

This paper presented a human-in-the-loop LLM-assisted multi-agent framework for distribution grid reliability planning. The proposed approach integrates LLM-based reasoning with deterministic reliability simulation tools through a coordinated multi-agent architecture. By combining data ingestion, planning optimization, reliability evaluation, and compliance verification within a unified workflow, the framework enables partially automated reliability planning while preserving engineering validation through simulation and human oversight.

Case studies on 100 real-world distribution feeders demonstrate that the proposed framework can significantly reduce planning time and manual engineering effort while maintaining acceptable planning accuracy for large-scale reliability screening studies. The results also highlight improvements in data processing efficiency, reliability optimization, and automated report generation.

Although manual engineering workflows still achieve the highest planning accuracy, the proposed LLM-assisted framework provides a scalable tool to support engineering decision-making in complex distribution systems.

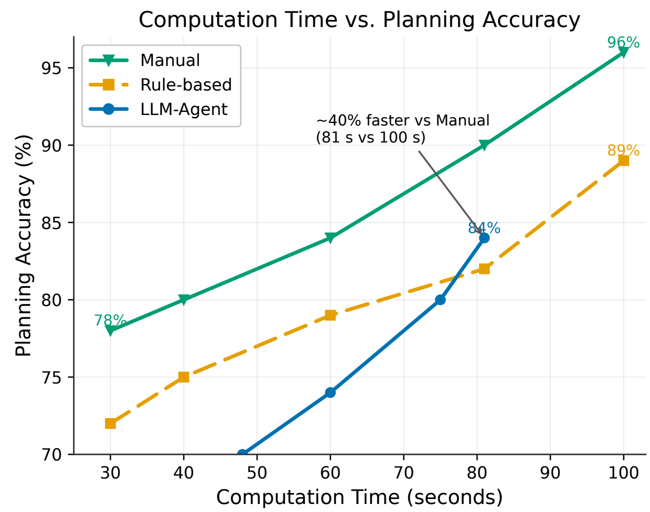

Fig. 4: Comparison of manual, rule-based, and LLM-agent workflows in terms of planning accuracy (A_p) and computation time.

Future work will focus on improving the integration between LLM reasoning and power system simulation tools, extending the framework to larger distribution networks, and exploring real-time reliability planning applications using streaming operational data.

References

[1] A. Rastgou, "Distribution network expansion planning: An updated review of current methods and new challenges," *Renewable and Sustainable Energy Reviews*, vol. 189, Art. no. 114062, Jan. 2024.

[2] M. Jooshaki *et al.*, "An enhanced MILP model for multistage reliability-constrained distribution network expansion planning," *IEEE Transactions on Power Systems*, vol. 37, no. 1, pp. 118–131, Jan. 2022.

[3] Z. Li *et al.*, "A reliability-constrained expansion planning model for mesh distribution networks," *IEEE Transactions on Power Systems*, vol. 36, no. 2, pp. 948–960, Mar. 2021.

[4] R. S. Pinto *et al.*, "Reliability-constrained robust expansion planning of active distribution networks," *IET Generation, Transmission & Distribution*, vol. 16, no. 1, pp. 27–40, Jan. 2022.

[5] Y. Zhang *et al.*, "Grid-Agent: An LLM-powered multi-agent system for power grid control," arXiv preprint arXiv:2508.05702, Sep. 2025.

[6] S. Choi *et al.*, "Generative AI for power grid operations," National Renewable Energy Laboratory (NREL), Golden, CO, USA, Tech. Rep. NREL/TP-5D00-91176, Nov. 2024.

[7] Y. Xiang *et al.*, "Large language model for secure operation of power systems," *Energy Systems Integration*, vol. 2, no. 1, pp. 1–15, 2025.

[8] R. S. Bonadia *et al.*, "On the potential of ChatGPT to generate distribution systems for load flow studies using OpenDSS," *IEEE Transactions on Power Systems*, early access, 2023.

[9] Z. Yan and Y. Xu, "Real-time optimal power flow with linguistic stipulations: Integrating GPT-agent and deep reinforcement learning," *IEEE Transactions on Power Systems*, vol. 39, no. 2, pp. 4747–4750, Mar. 2024.

[10] M. Jia, Z. Cui, and G. Hug, "Enhancing LLMs for power system simulations: A feedback-driven multi-agent framework," *IEEE Transactions on Smart Grid*, vol. 16, no. 6, pp. 5556–5572, Nov. 2025.

[11] B. Shen *et al.*, "Large language model-based security situation awareness for smart grid: Framework and approaches," *IEEE Access*, vol. 13, pp. 173600–173613, 2025.

979-8-3315-8138-1