A Locally-Hosted Agentic LLM Framework for Resilience-Focused Power System Planning

Eric Keller, Jonathan M. Snodgrass, Thomas J. Overbye
Department of Electrical & Computer Engineering
Texas A&M University
College Station, TX, USA
{erickeller, snodgrass, overbye}@tamu.edu

***Abstract*—Large language models (LLMs) are increasingly capable tool-using agents, yet their application to power-system planning is constrained by data-confidentiality, model-size, and reliability requirements that cloud-hosted models do not meet. This paper proposes a locally hosted, tool-calling agentic architecture for resilience-focused planning studies. The system couples a central orchestrator, a retrieval-augmented-generation (RAG) layer over regulatory documents and historical cases, a hierarchy of scripts, tools, and skills that drive power flow modeling and simulation sowftware, and an interpretive layer that presents results to a user who retains final approval of any grid action. Because hazardous operating conditions are rare, a domain-reweighting scheme adapted from DoReMi balances normal and hazardous data across multiple hazard severities. Our initial implementation is tested on a lightweight mixture-of-experts model on a desktop unit, with a much larger model on more powerful, gpu-equipped server to ensure that the larger model can benefit from efficiencies needed to aid the smaller one, and it demonstrates fragility-based outage-scenario generation for tornadoes (1950–present) and hurricanes on a synthetic, roughly 2000-bus Texas case. Grounded, locally hosted agents can substantially reduce the setup and execution burden of hazard-aware planning studies.**

***Index Terms*—Power Systems, LLMs, agentic AI, grid resilience, transmission planning, RAG, hazard modeling**

I. Introduction

Large language models (LLMs) have rapidly become more capable and more widely used. Originally text-to-text systems, modern models add data retrieval, image generation and interpretation, and tool use, and they increasingly operate as agents that decompose a request into actions [1]. Power systems, in turn, produce large volumes of data in many formats, and an exploration of how LLMs can assist in interpreting that data is in order. It is our belief that LLMs show promise in any text-related facet of the field—regulatory reference, coding against power-simulation software, and cybersecurity—but that they require a supportive architecture to meet the field's standards for safe operation.

The combined field is, however, difficult to survey. The first issue is scope. Some approaches operate at the operational level [2], others at the device level [3], and others, as we do, in the planning and study space. Instead of grappling with the entirety of the extant literature on LLMs as they are structured and deployed in what similar contexts as we could find, opted for this proposal rather than a full review. We reviewed solutions to expected or encountered challenges in our design, thereby limiting our scope to that which is directly useful.

That scope is threefold: (1) to build an agentic planning and study system that (2) operates locally, on protected data, and (3) provides general power-expert interaction to its users. Automating the creation and execution of studies—even without the analysis—would save grid planners time. Local operation is required because real case information is controlled. Field-specific interaction reduces friction for the expected user base, who are informed in power systems if not always experts, by "speaking their language" in a technical rather than a cultural sense.

The contributions of this paper are:

1) a locally hosted, tool-calling agentic architecture—an orchestrator, a RAG layer, a script/tool/skill hierarchy that drives powerflow simulation and grid modeling software, and an interpretive layer—that runs on controlled data without cloud dependence while keeping the human planner in the approval loop;
2) a resilience-oriented data-balancing approach that adapts DoReMi-style domain reweighting to mix normal and hazardous conditions across hazard types and severities; and
3) an initial implementation that demonstrates fragility-based outage-scenario generation for tornadoes and hurricanes on a synthetic, roughly 2000-bus Texas case.

The remainder of the paper is organized as follows. Section II reviews related work. Section III presents the proposed system design. Section IV discusses the principal design challenges and our mitigations. Section V reports implementation status and planned work, and Section VI concludes.

II. Related Work

LLM applications to power systems have grown quickly; a recent review catalogs work spanning fault diagnosis, forecasting, control, and planning [1]. We organize the most relevant

Funding was provided by the Department of Energy (DOE) Office of Cybersecurity, Energy Security, and Emergency Response (CESER) through project DE-CR0000067. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Copyright ©2026 IEEE. Personal use of this material is permitted. North American Power Symposium (NAPS) 2026, Houghton, MI, October 2026.

efforts into agentic systems for grid analysis and control, and LLMs for forecasting and cross-domain reasoning.

A. Agentic Systems for Grid Analysis and Control

The closest comparison to our work is the agentic static-analysis tool of Wang *et al.* [4], which converts natural-language requests into MATPOWER scripts and adds a three-tier error-correction system—a static pre-check, a dynamic feedback loop, and a semantic validator—reporting roughly 82% code fidelity. Our intent is to use a unified package that targets a variety of power-modeling software, initially PowerWorld Simulator, so that the coding agent works against a simplified, curated interface; we expect this to reduce the number of errors and lessen the need for an extensive error-correction suite. Like that work, we make use of the model context protocol (MCP) [5].

Several systems place an LLM in or around a control loop. Zhang *et al.* [2] integrate an LLM with a safety-constrained soft actor–critic for topology reconfiguration, using the LLM to steer the agent away from unsafe switching actions on the IEEE 36- and 118-bus Grid2Op benchmarks. Chehade and Zhu [6] fine-tune an instruction-tuned LLM to produce economical, reliable corrective switching plans, reducing AC power-flow failures from 50% to single digits on IEEE 118-bus public-safety-power-shutoff scenarios. Yang *et al.* [3] pair an LLM with reinforcement learning for two-stage distribution-network voltage control. A recurring concern across these efforts is scale: test grids of 118 buses are common, whereas planning-scale grids contain thousands of buses, which raises a real question of feasibility.

For optimization specifically, PowerGraph-LLM [7] couples graph and tabular grid representations to query an LLM for optimal power flow (OPF). The limits of direct LLM optimization are equally instructive: Bernier *et al.* [8] construct an OPF benchmark and find that state-of-the-art and reasoning models alike fail on the harder constrained settings, which reinforces our decision to keep numerical work inside validated solvers rather than the model. Multi-agent designs such as Grid-Agent [9] and the hierarchical X-GridAgent [10] use tool-calling and RAG to perform contingency and broader steady-state analysis through natural-language queries; the interpretive layer of our system builds on this line of work [10].

B. LLMs for Forecasting and Cross-Domain Reasoning

Much of our data is time-variant—weather, prices, and demand—and LLMs have been adapted to forecast such series: TimeCMA [11] aligns time-series and text representations for multivariate forecasting. The grid is not the only network amenable to this treatment. Like power systems, road and cellular traffic networks are naturally modeled with nodes and links; Yang *et al.* [12] forecast spatiotemporal traffic by feeding traffic matrices to a vision-capable LLM as images. As traffic flows vary over time and links experience congestion, so power flows shift with prices, weather, and load response, and lines experience overloads rather than travel-time increases—an analogy that suggests visualization tools embedded in power-modeling software could be adapted for visual processing. Finally, because modern grids include digital sensing and control whose output is renderable as text, monitoring methods transfer as well: Worae and Mastorakis [13] interpret IoT packet traffic for anomalies, and supervisory control and data acquisition (SCADA) telemetry differs from IoT traffic in type but not in form.

C. Gaps and Differentiation

Grid control [9], optimal power flow [7], and general analysis [10] are already represented in the literature. What remains underexplored is an agentic system that rapidly creates, modifies, and analyzes grids specifically in the context of reliability and resilience to hazards for planning purposes. This is the gap we address, and we argue that LLM agents are well suited to assist here as well.

III. Proposed System Design

We propose a three-part system, two parts of which are covered here; the third, an interpretive layer, builds on the analysis agent of Wen and Chen [10]. As shown in Fig. 1, an agent sits at the center of the system with access to static data and records through RAG and to an extensive tool set with which it manipulates power-simulation software. The software produces simulation results, which are passed to the interpretive layer for presentation to the user. The end goal is an all-in-one agentic assistant that carries the background, training, and data needed to off-load previously tedious setup, execution, and reference tasks, allowing users to expand the scope of their studies. The same package is intended to help forecast grid health from publicly available data and to supply the study data needed to train that forecasting and interpretive function, alongside existing weather-into-power-flow studies [14]–[16]. For grids, we draw on the synthetic networks of the Texas A&M repository [17] and on workflows from prior group tools such as the EIA-860 case-creation pipeline [18], [19]. Fig. 2 details the runtime architecture, from startup and chat interface through the agentic/non-agentic branch to the tool and modeling layers. Final approval rests with the user, who is ultimately responsible for any action taken with respect to the grid, whether direct or indirect.

A. Power Simulation Package

The system is currently tied to a single software solution, although we have no opposition to adding support for other packages where possible. We add the caveat that the available tools are bounded by the feature set of the installed software. The package is organized into three layers. *Scripts* have no fixed scope: in principle, every function exposed by the simulator's application programming interface (API) can be collected into a single function file that tool scripts then call. Our repository incorporates scripts written by past group researchers, which—because they represent real workflows by real users—became the basis for the higher layers. *Tools* are more itemized: the ideal tool is a single script that accomplishes a single task, or a small bundle of tasks that must

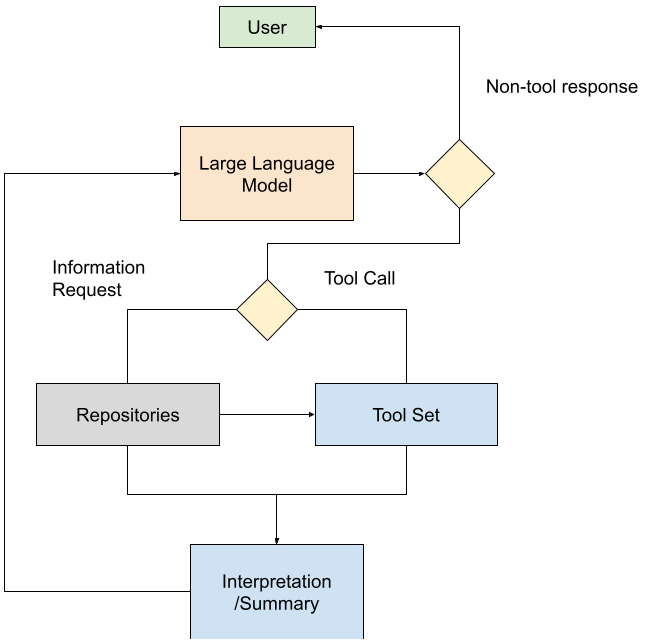

Fig. 1. System high-level information flow. The agent either returns a non-tool response to the user or issues a tool call; tool calls retrieve data from repositories or act through the tool set, and the resulting interpretation is returned to the agent.

be grouped, such as opening the simulator and a case. *Skills* are complete workflows assembled from tools; for example, a weather study may open the case, insert power-flow snapshots, and then select and run the desired results as three tools bundled into one skill. Because one user's complete workflow is another's partial workflow, skills can overlap and are the natural place for user customization: a study of a single tornado and a batched study of all tornadoes over a time window share most of the same steps.

IV. Design Challenges and Mitigations

Realizing this system raises several challenges, addressed in turn.

A. Locality and Model Selection

A core requirement is that the LLM be hosted and operated locally, so that the agent can work with controlled data without pushing it across unsecured channels or to uncleared parties, including the frontier laboratories that provide today's most capable models. Models can be obtained and run through repositories such as Hugging Face [20]. Because the system relies on a tool package, it must use—or train—a tool-calling model; for convenience we adopt the former, using gpt-oss-120B for deployment, with package testing performed on the smaller mixture-of-experts (MoE) model gemma-4-a4b-26b. Locality imposes a de facto size limit: several models must coexist in the same infrastructure, and the most capable cloud checkpoints would be too large for consumer- or prosumer-grade hardware, so an isolated deployment of those models would require building or securing a data center. The compact MoE model is favored precisely because its roughly 12 GB footprint fits a single workstation.

B. Grounding via MCP and RAG

The model context protocol is Anthropic's standard for granting a model access to files and functions in a host environment [5]. Several MCP "servers," locally hosted, are surfaced to the model, which is itself loaded on a local server; this yields an agent with access to real documents and scripts, which we expect to reduce the fine-tuning and low-rank-adaptation effort otherwise needed. Rather than being trained to answer questions about a standard, the agent can look up the standard itself. Retrieval-augmented generation (RAG) serves the same end [21], exposing facts for the model to report or summarize rather than relying on its parameters to produce guesses. The data exposed include regulatory documents, cases, study results, and historical hazard records. This retrieval store is the heaviest part of the system; an end user without sufficient memory or access would retrieve and remove only what a given study needs, retaining assistance—if less than a fully provisioned database would offer.

C. Reliability and Context Efficiency

LLMs tend to present a fluent but unsupported account of a topic rather than the observable truth, a behavior commonly termed hallucination [22]. We do not propose a complete solution here, nor do we investigate causes; mitigation methods exist, including the use of additional agents, and future work may incorporate corrective systems. Our focus is instead on supplying the data and the grounding needed to make hallucination less likely. A related, practical issue is context efficiency. In testing, the local model would expend so much of its budget relating its reasoning to the user that it failed to produce an answer; context-window-extension techniques such as positional-interpolation methods [23], or modifications like adding attention layers to expand context [24], together with a summarization stage relieve this, and a summarization process is in any case required because the full size of some cases and result sets exceeds what the model can ingest directly. Where domain adaptation is unavoidable—proprietary simulator APIs have little footprint in public models—we plan to adapt models through low-rank adaptation (LoRA) [25] and fine-tuning [26], training on API documentation and on code examples drawn from our own and public code bases. Pre-organization of code into example scripts and workflows further limit the need for on-the-fly coding, reducing the work burden of the user. Text-to-text components are judged by BERTScore [27], and code is judged by its ability to execute specific tasks.

D. Data Imbalance

The focus of this work is protection against, management of, and recovery from hazardous conditions, which are by their nature unusual. A model trained to forecast and monitor a grid would normally be trained on a stable grid and would therefore learn to predict stability whatever happens [28]. Hazards also occur at different severities, which are used

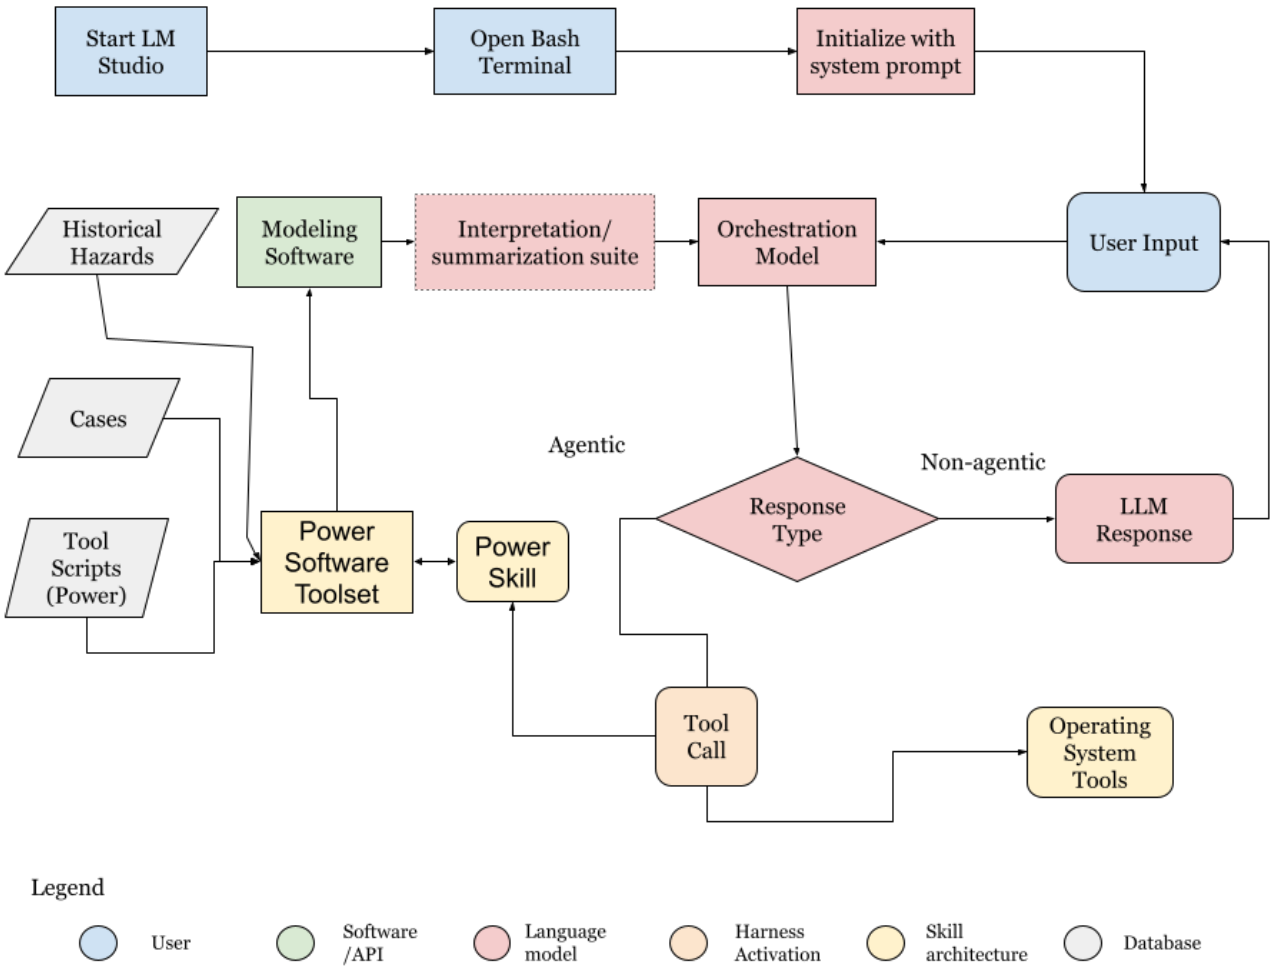

Fig. 2. The expanded information flow. The database, tool, and request handling are expanded.

to gauge potential damage: tornadoes are rated on the six-level Enhanced Fujita scale [29] and hurricanes on the five-category Saffir–Simpson wind scale [30]. These severities are themselves sub-categories, so a useful mixture is needed both within categories and between normal and hazardous operation. We adapt DoReMi-style domain reweighting [31] to this setting: where the original method sets domains by data source, ours sets them by operating regime—normal versus hazardous, further divided by hazard type and severity. Concretely, given K domains, domain weights α on the simplex Δ^K are optimized to minimize the worst-case excess loss relative to a reference model,

$$\min_{\theta} \max_{\alpha \in \Delta^K} \sum_{d=1}^{K} \alpha_d \big[\ell_d(\theta) - \ell_d(\theta_{\text{ref}}) \big], \tag{1}$$

where $\ell_d(\theta)$ is the loss of model θ on domain d; the resulting weights then guide resampling of the training data. We prefer automated balancing over manual labeling, and prior group work measuring the number and severity of renewable-energy droughts [15] and geomagnetic disturbances [32], [33] supplies labeled hazardous data to draw on.

V. Implementation Status and Future Work

Most basic power-modeling functions have been incorporated into the tool set, and several workflows—including the weather-scenario process and case-creation functions [18], [19]—have been transformed into tools and skills. Historical tornado data from 1950 to the present has been assembled. To estimate grid impact, the tool set includes scripts that handle single cyclonic events and tornado clusters arising from supercells. Severity, as measured on the Enhanced Fujita scale, is used to construct a wind field whose intensity is compared against grid-object fragility curves. For an asset i at time step t, let $v_{i,t}$ be the local wind speed and $F_i(\cdot)$ its fragility curve, giving the probability of failure as a function of wind speed. The asset is taken offline when a random draw $u_{i,t} \sim \mathcal{U}(0,1)$ satisfies

$$u_{i,t} \le F_i(v_{i,t}) \,, \tag{2}$$

and a time-stamped contingency is then created for that asset. The wind field is advanced to the tornado's next position—estimated or actual—at the following time step, and the evaluation is repeated. A similar method drives the hurricane generator, whose historical track data, available from sources such as the NOAA Atlantic hurricane database [34], has not yet been assembled.

For the tornado case, a batch of historical runs was performed on our synthetic, roughly 2000-bus Texas case, comprising scenarios for each individual tornado and each identified cluster from 1950 through 1959. The hurricane generator was exercised by creating an arbitrary Category 3 storm and routing its track through the Houston and Dallas areas. In both cases the code is not yet finalized; the runs confirmed that the process produces the intended output—a failure frequency at each time step—but further testing is required before the results can be trusted.

Several extensions are planned. The wind fields will be revised to use minimum, maximum, and median values for the Enhanced Fujita and Saffir–Simpson categories when setting field strength, with real data used where available, which is rarely possible for tornadoes. Other groups are developing additional hazard scenarios that will be incorporated when available; in particular, previously developed code for geomagnetic-disturbance scenarios, together with results already obtained on several of our grids [32], [33], will be folded into the package and used to identify worst-case events, which will speed hazard selection for training. Because hazardous conditions span event types that each vary in severity, multiple layers of data balancing may ultimately be required.

VI. Conclusion

There is rough consensus on the components a grid-planning assistant should include. We differentiate this work by focusing on disaster scenarios and by incorporating a more complete, more thoroughly grounded agent workflow. We do not intend to replace human engineers; on the contrary, we believe human oversight is essential to the safe use of a tool like this. The aim is to empower planners and, through the data produced in its construction, to benefit future researchers and industry practitioners alike.

References

[1] F. Amjad, T. Korõtko, and A. Rosin, "Review of LLMs applications in electrical power and energy systems," *IEEE Access*, vol. 13, pp. 150 951–150 969, 2025.

[2] Z. Zhang, C. Shen, X. Wan, J. Song, and M. Sun, "LLM-guided safe reinforcement learning for energy system topology reconfiguration," arXiv preprint arXiv:2603.14018, 2026.

[3] X. Yang, C. Lin, X. Ma, D. Liu, R. Zheng, H. Liu, and W. Wu, "Two-stage active distribution network voltage control via LLM-RL collaboration: A hybrid knowledge-data-driven approach," arXiv preprint arXiv:2602.21715, 2026.

[4] Q. Wang, S. Yang, and Y. Zhu, "Agentic application in power grid static analysis: Automatic code generation and error correction," arXiv preprint arXiv:2604.09995, 2026.

[5] Anthropic, "Introducing the model context protocol," https://www.anthropic.com/news/model-context-protocol, 2024, accessed: 2026.

[6] M. Chehade and H. Zhu, "Fine-tuning LLMs to generate economical and reliable actions for the power grid," arXiv preprint arXiv:2602.15350, 2026.

[7] F. Bernier, J. Cao, M. Cordy, and S. Ghamizi, "PowerGraph-LLM: Novel power grid graph embedding and optimization with large language models," *IEEE Transactions on Power Systems*, vol. 40, no. 6, pp. 5483–5486, 2025.

[8] F. Bernier, S. Ghamizi, P. Dogoulis, and M. Cordy, "Can large language models reason and optimize under constraints?" arXiv preprint arXiv:2603.23004, 2026.

[9] Y. Zhang, A. M. Saber, A. Youssef, and D. Kundur, "Grid-Agent: An LLM-powered multi-agent system for power grid control," arXiv preprint arXiv:2508.05702, 2025.

[10] Y. Wen and X. Chen, "X-GridAgent: An LLM-powered agentic AI system for assisting power grid analysis," arXiv preprint arXiv:2512.20789, 2025.

[11] C. Liu, Q. Xu, H. Miao, S. Yang, L. Zhang, C. Long, Z. Li, and R. Zhao, "TimeCMA: Towards LLM-empowered multivariate time series forecasting via cross-modality alignment," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2025.

[12] N. Yang, H. Zhong, H. Zhang, and R. Berry, "Vision-LLMs for spatiotemporal traffic forecasting," arXiv preprint arXiv:2510.11282, 2025.

[13] D. A. Worae and S. Mastorakis, "An LLM-powered AI agent framework for holistic IoT traffic interpretation," arXiv preprint arXiv:2510.13925, 2025.

[14] F. Safdarian, J. S. Cook, K. Zhgun, T. J. Overbye, and J. Snodgrass, "Power flow modeling of the impacts of weather and other resiliency hazards with a focus on transmission planning," in *Proc. 58th Hawaii International Conference on System Sciences (HICSS)*, Waikoloa, HI, USA, Jan. 2025.

[15] K. Zhgun, F. Safdarian, and T. J. Overbye, "Classification and identification of renewable energy droughts in the united states," in *2025 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, USA, Feb. 2025.

[16] S. J. Lee, F. Safdarian, and T. J. Overbye, "Validation of wind and PV power generation using historical and forecast weather data," in *2025 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, USA, Feb. 2025.

[17] S. Kunkolienkar, F. Safdarian, J. Snodgrass, A. B. Birchfield, and T. J. Overbye, "A description of the Texas A&M university electric grid test case repository for power system studies," in *2024 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, USA, Feb. 2024.

[18] E. Keller, W. Yun, J. M. Snodgrass, and T. J. Overbye, "CaseBuilder and the EIA-860 dataset: Towards a rapid power-case creation and modification system," in *2026 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, USA, Feb. 2026. [Online]. Available: https://snodgrass.engr.tamu.edu/wp-content/uploads/sites/350/2026/01/CaseBuilder_TPEC2026_Website.pdf

[19] J. S. Cook, F. Safdarian, J. Snodgrass, and T. J. Overbye, "Using power flow application capabilities to visualize and analyze US energy information administration generation data," in *2024 IEEE Texas Power and Energy Conference (TPEC)*, College Station, TX, USA, Feb. 2024. [Online]. Available: https://overbye.engr.tamu.edu/wp-content/uploads/sites/146/2024/01/EIA-Cookbok-Jan-4_ARCHIVE.pdf

[20] Hugging Face, "Hugging face: The AI community building the future," https://huggingface.co/, 2026, accessed: 2026.

[21] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020.

[22] J.-Y. Yao, K.-P. Ning, Z.-H. Liu, M.-N. Ning, and L. Yuan, "LLM lies: Hallucinations are not bugs, but features as adversarial examples," arXiv preprint arXiv:2310.01469, 2023.

[23] B. Peng, J. Quesnelle, H. Fan, and E. Shippole, "YaRN: Efficient context window extension of large language models," in *International Conference on Learning Representations (ICLR)*, 2024.

[24] H. Jin, X. Han, J. Yang, Z. Jiang, Z. Liu, C.-Y. Chang, H. Chen, and X. Hu, "Llm maybe longlm: Self-extend llm context window without tuning," arXiv.org, 02 2024. [Online]. Available: https://arxiv.org/abs/2401.01325

[25] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *International Conference on Learning Representations (ICLR)*, 2022.

[26] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, "Finetuned language models are zero-shot learners," in *International Conference on Learning Representations (ICLR)*, 2022.

[27] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "BERTScore: Evaluating text generation with BERT," in *International Conference on Learning Representations (ICLR)*, 2020.

[28] M. Abdulrazaq, "Rare-event prediction in imbalanced data: A unified evaluation and optimization framework for high-risk systems," *Communication in Physical Sciences*, vol. 9, no. 4, pp. 968–979, 2023.

[29] National Weather Service, "The enhanced Fujita scale (EF scale)," https://www.weather.gov/oun/efscale, 2019, accessed: 2026.

[30] National Hurricane Center, "Saffir-Simpson hurricane wind scale," https://www.nhc.noaa.gov/aboutsshws.php, 2025, accessed: 2026.

[31] S. M. Xie, H. Pham, X. Dong, N. Du, H. Liu, Y. Lu, P. Liang, Q. V. Le, T. Ma, and A. W. Yu, "DoReMi: Optimizing data mixtures speeds up language model pretraining," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 36, 2023.

[32] P. Dehghanian, A. Zhang, R. Fatima, J. Snodgrass, A. B. Birchfield, K. R. Davis, and T. J. Overbye, "An integrated assessment of a G3 GMD event on large-scale power grids: From magnetometer data to geomagnetically induced current analysis," *IEEE Transactions on Industry Applications*, vol. 60, no. 1, pp. 1634–1644, 2024. [Online]. Available: https://overbye.engr.tamu.edu/wp-content/uploads/sites/146/2024/08/IAS_Paper___TPEC_Extension_WithFooter_2024.pdf

[33] N. LoGuidice, J. Snodgrass, and T. J. Overbye, "Estimating the electric field from geomagnetically induced currents," in *2024 IEEE Kansas Power and Energy Conference (KPEC)*, Manhattan, KS, USA, Apr. 2024. [Online]. Available: https://overbye.engr.tamu.edu/wp-content/uploads/sites/146/2024/04/E_Field_State_Estimator_Using_GIC_Measurements_Digest_Footer_copyright_ARCHIVE.pdf

[34] National Oceanic and Atmospheric Administration, "NOAA atlantic hurricane database (HURDAT2)," https://www.kaggle.com/datasets/utkarshx27/noaa-atlantic-hurricane-database, 2024, accessed: 2026.