

Machine Learning for Predictive Cost and Schedule Risk Mitigation in Inter-State Transmission Line Buildouts

Lukhanyiso Mpongoshé,^{1,*} and Chidinma Queen Adieze²

1: Independent Infrastructure Project Delivery Consultant, Midrand, Gauteng, South Africa

2: Department of Data Analytics / Western Governors University, Salt Lake City, UT, United States

Received July 17, 2026; Accepted August 24, 2026; Published September 3, 2026

Inter-state transmission projects in the United States routinely exceed their approved schedules and budgets, and the point of maximum financial exposure, sanction, when capital is irreversibly committed, is typically reached before the largest sources of delay are quantified. This paper proposes a machine-learning framework for forecasting cost and schedule risk on inter-state transmission line buildouts at the pre-sanction stage, using only information available before financial commitment. The framework draws on publicly reported data layers, including transmission plans filed under FERC Form 715, reliability and equipment-addition records from NERC and the U.S. Energy Information Administration, environmental permitting timelines, interconnection-queue position from the Lawrence Berkeley National Laboratory Queued Up series, and published equipment lead-time indices, and it combines a gradient-boosted ensemble for overrun magnitude with a survival model for milestone timing. We specify the data architecture, feature design, target definitions, and an evaluation protocol appropriate for a domain with few projects and long horizons, and we discuss governance and interpretability requirements for adoption by utilities, transmission developers, and their engineering, procurement, and construction partners. The paper is a framework and methodology contribution: illustrative figures denote expected behaviour and are not empirical results. We conclude with an implementation pathway and the data-access steps required for empirical validation.

Keywords: Transmission Planning; Project Risk; Machine Learning; Cost Overrun; Schedule Delay; Supply Chain

Introduction

The United States is entering a period of transmission expansion without a recent precedent. The Department of Energy's National Transmission Planning Study concluded that meeting demand affordably and reliably will require a substantial enlargement of the transmission system, potentially doubling its extent by mid-century, and federal programmes have been directed toward accelerating reconductoring and advanced transmission technologies. At the same time, load forecasts have turned sharply upward after years of flat growth, driven by data-center and industrial demand. The gap between what must be built and what the delivery system can absorb is now the binding constraint on modernization.

*Corresponding author: lukhanyisompongoshé@gmail.com

That constraint is increasingly a delivery problem rather than a financing or engineering problem. Lead times for large power transformers, historically between seven and fourteen months, have extended in many cases beyond two years, and for specialized high-capacity units toward three to four years. Because a transmission line cannot energize without its terminal equipment, procurement timing has become a first-order determinant of whether a project meets its in-service date. Yet in conventional practice, the decision that commits capital, project sanction, is taken before these delay drivers are quantified with any rigor. The result is a systematic asymmetry: risk is realized late but committed early.

This paper asks whether the information needed to anticipate cost and schedule risk already exists in the public record before sanction, and whether machine-learning methods can convert that information into calibrated, decision-useful forecasts. Inter-state lines are the focus because they concentrate the hardest risks (multiple regulatory jurisdictions, longer routes, contested rights-of-way, and larger equipment packages) and because they are precisely the projects the national planning agenda depends upon. The contribution is a framework: a specification of data sources, features, model structure, and evaluation protocol, together with the governance requirements for real-world adoption. It is not an empirical study; no model is trained here, and all figures are illustrative. The final section outlines exactly which data access is required to move from specification to validation.

The remainder of the paper is organized as follows. The next section reviews prior work on infrastructure cost and schedule forecasting and situates the present approach. The Materials and Methods section specifies the data layers, feature engineering, model architecture, and evaluation design. The Results and Discussion section presents the framework's expected behaviour through illustrative figures and discusses interpretability, governance, and adoption. A short Conclusions section summarizes the implementation pathway.

Background and Related Work

Three strands of prior work bear on the problem addressed here: the long-standing study of cost and schedule performance in large infrastructure, the more recent application of machine learning to construction and energy-project prediction, and the emerging literature on transmission supply-chain constraints. We review each in turn and then state the gap that motivates the present framework.

Cost and Schedule Performance of Large Transmission Projects

Large linear infrastructure projects have long exhibited optimism bias and strategic misrepresentation in early estimates, a pattern documented across transport and energy megaprojects [5, 6]. Reference-class forecasting was proposed to counter this by anchoring estimates in the distribution of outcomes from comparable completed projects rather than in bottom-up build-ups alone [5]. The approach is powerful but coarse: it produces a single distribution for a reference class and does not condition on the specific, changing risk drivers of an individual project. Recent syntheses continue to find that transmission and grid projects are prone to escalation, and that the causes are dominated by permitting, routing, and stakeholder factors that are, in principle, visible well before construction begins [13, 14]. This is important for the present work because it implies

that the information needed to anticipate escalation is largely present at the pre-sanction stage rather than emerging only during delivery. Machine-learning methods extend the reference-class logic by learning, from many projects, how outcomes vary with observable project characteristics, thereby replacing a single class-level distribution with a conditional, project-specific one.

Machine Learning for Construction and Energy-Project Prediction

Predictive analytics are now routine in the operation of power systems, including load forecasting, predictive maintenance, and outage reduction [11]. A rapidly growing body of work applies supervised learning to construction cost estimation and schedule prediction, with tree-based ensembles such as gradient boosting repeatedly reported as strong performers on the tabular, mixed-type data typical of project records [7, 9]. Reviews published between 2023 and 2026 report expanding use of ensemble and hybrid models for construction cost and duration prediction, alongside a consistent set of unresolved problems: small and heterogeneous datasets, weak external validation, and limited model transparency [15, 16, 18]. The transparency problem is particularly acute for capital-committee use, where an unexplained score will not be trusted; post-hoc attribution methods such as additive feature explanations have therefore become a standard companion to ensemble models [10, 19]. Two limitations recur across this literature and shape our design. First, most studies predict outcomes using information that becomes available only during or after execution, which limits their value for a go or no-go decision. Second, few studies treat equipment procurement lead time as a predictor at all, reflecting the fact that much of the evidence base predates the current supply environment. The framework proposed here is deliberately restricted to pre-sanction information and elevates procurement lead time to a first-class feature, accepting weaker inputs in exchange for decision relevance.

The Transmission Supply-Chain Dimension

Supply conditions since 2021 have made equipment lead time a dominant schedule driver rather than a background planning assumption. Independent assessments report that lead times for large power transformers have roughly doubled since 2021, extending beyond two years in most major markets and toward three to four years for specialized high-capacity units, with record manufacturer backlogs [8, 17]. Demand for generator step-up transformers rose several-fold over 2019 to 2025 and prices increased sharply over the same period, driven by data-centre load, electrification, and grid modernization arriving simultaneously [20, 21]. Grain-oriented electrical steel, essential to transformer cores, is domestically produced at only a single U.S. source, concentrating supply risk [21]. Analysts differ on whether the binding constraint is manufacturing capacity or procurement practice, but they agree that procurement timing now determines which projects advance and which stall [17, 20]. Any credible pre-sanction risk model must therefore explicitly represent procurement exposure and revalidate it as supply conditions change, since a lead-time relationship learned in one procurement environment will not transfer to another.

Gap and Contribution

Taken together, the literature establishes that transmission cost and schedule escalation is common, that its principal drivers are observable before sanction, and that machine learning can produce conditional risk estimates from project-level data, yet no

established framework combines these observations into a pre-sanction, procurement-aware, interpretable risk tool specific to U.S. inter-state transmission. The present paper addresses that gap at the level of specification: it defines the data architecture, features, targets, model structure, evaluation protocol, and required governance conditions, and sets out the data-access pathway for empirical validation.

Materials and Methods

The framework is described at the level of data, features, targets, model, and evaluation. Because the paper is a methodological specification, this section is written so that a researcher with access to the data can implement and test it directly.

Data Layers

Five public or semi-public data layers supply the model. Each is available, in principle, before project sanction, which is the defining constraint of the design.

Table 1. Public data layers and their role in the framework

Data layer	Representative source	Role in model
Transmission plans	FERC Form 715 annual filings	Project scope, voltage, route, and planned in-service date
Reliability & additions	NERC; EIA Form 411 / 861	Historical additions, utility delivery history
Permitting timelines	NEPA records; state PUC / siting dockets	Environmental and siting duration, jurisdictions crossed
Equipment lead times	Published transformer & conductor indices	Procurement exposure to sanctions
Interconnection queue	LBNL 'Queued Up' dataset	Queue position, cluster study status, withdrawal risk

Feature Engineering

Raw records are aligned to a common project key and converted into features observable at sanction. These fall into four groups: **physical** (route length, voltage class, terminal count, corridor terrain); **jurisdictional** (number of states and permitting authorities crossed, presence of contested right-of-way, siting-docket age); **procurement** (transformer and conductor lead-time indices at the planned order date, degree of single-source exposure); and **institutional** (the sponsoring utility's historical variance between planned and actual in-service dates, derived from prior filings). The institutional group is what allows the model to learn that identical projects carry different risk depending on who delivers them.

Target Variables

Two targets are modeled. The first is **cost overrun**, defined as the ratio of final to sanction-stage approved cost, treated as a continuous regression target and reported as a distribution rather than a point value. The second is **schedule slippage**, defined as the difference between realized and planned in-service dates, modelled per milestone so that the framework can indicate not only whether a project will be late but at which stage the delay is most likely to originate.

Model Architecture

The core combines two complementary learners. A gradient-boosted regression ensemble estimates the overrun magnitude and, through quantile objectives, its uncertainty band. A discrete-time survival model estimates the probability that each milestone is reached by its planned date, which suits censored data (projects still in progress) better than ordinary regression. Figure 1 shows the overall pipeline from the public data layers through feature alignment to the decision outputs.




Figure 1. Conceptual architecture of the pre-sanction risk framework, from public data layers (left) through feature engineering to a gradient-boosted ensemble and survival model, producing milestone delay probabilities, a cost-overrun distribution, and a risk-ranked portfolio view.

Evaluation Design

Transmission projects are few and slow, so a naive random train-test split would leak information across time. Evaluation therefore uses forward-chaining validation: the model is trained on projects sanctioned up to a cut-off year and tested on those sanctioned afterward, repeated across rolling cut-offs. Calibration, meaning whether a stated eighty-percent confidence band contains the outcome eighty percent of the time, is treated as more important than point accuracy, because the intended use is risk communication at a go/no-go decision, not precise cost prediction.

Results and Discussion

This section presents the framework's expected behaviour. **The figures in this section are illustrative and synthetic; they depict how the framework is intended to perform and are not results from trained models.** They are included, in keeping with the journal's guidance on clear presentation, to make the proposed behaviour concrete.

The Delivery Constraint in Context

Figure 2 places the equipment-lead-time problem on a timeline. The reported expansion of large-transformer lead times, from a pre-2020 norm of roughly seven to fourteen months toward two to four years, is the single change that most alters pre-sanction risk, because it moves procurement from a manageable downstream activity to a

binding upstream constraint. A risk model that omits it would misstate schedule exposure for most current inter-state projects.




Figure 2. Illustrative depiction of reported large-power-transformer lead-time expansion. The band denotes the range reported across industry sources; the dashed line marks the pre-2020 norm. Values are indicative and shown for context, not as measured data.

Expected Calibration

Figure 3 shows the intended calibration behaviour on synthetic data: predicted schedule slippage tracks realized slippage along the diagonal, with dispersion that widens for longer-horizon projects. In deployment, the meaningful test is not that points lie exactly on the line but that the model's stated uncertainty bands are honest, a property the evaluation design explicitly targets.




Figure 3. Illustrative (synthetic) calibration of predicted versus realized schedule slippage. Generated for exposition only; not an empirical result.

Interpretability and the Decision Use

For a go/no-go decision, a risk score is useful only if its drivers are legible to the decision-maker. Figure 4 shows the kind of feature-contribution ranking the framework is designed to produce, with procurement lead time and permitting duration expected to dominate, followed by route and jurisdictional complexity. Such rankings translate directly into management action, such as accelerating a transformer order, sequencing permitting differently, or staging a right-of-way strategy, which is the point of forecasting before, rather than after, commitment.




Figure 4. Illustrative feature-contribution ranking. Ordering reflects domain expectation and is shown to convey the intended interpretive output; values are not estimated from data.

Governance, Limitations, and Adoption

A model of this kind should inform judgment, not replace it. Three governance conditions matter for adoption. First, the model must be transparent at the level of individual predictions, so that a capital committee can see why a project is flagged. Second, it must be revalidated as supply conditions change; a lead-time relationship learned in one procurement environment will not hold in another. Third, accountability for the sanction decision must remain with the responsible officers, with the model positioned as a structured second opinion. The principal limitations are data sparsity, since interstate projects are few, and reporting inconsistency across jurisdictions, both of which the evaluation design accommodates but cannot eliminate. These constraints argue for a consortium approach in which utilities, developers, and their engineering, procurement, and construction partners contribute delivery histories under a shared, governed data standard.

CONCLUSIONS

The binding constraint on U.S. inter-state transmission expansion has shifted from financing to delivery, and the delay drivers that most threaten delivery, namely

equipment lead times, permitting duration, and jurisdictional complexity, are largely observable before capital is committed. This paper specifies a machine-learning framework that converts pre-sanction information into calibrated cost and schedule risk forecasts by combining a gradient-boosted ensemble for overrun magnitude with a survival model for milestone timing, and sets out the data architecture, features, targets, evaluation protocol, and governance conditions for its use. The immediate next step is empirical: assembling a linked dataset from FERC Form 715 filings, NERC and EIA records, permitting dockets, published lead-time indices, and the LBNL interconnection-queue series, and validating the framework under forward-chaining evaluation. Realized as a governed, multi-party capability, such a tool would let decision-makers act on transmission delivery risk while the decision is still reversible.

ACKNOWLEDGMENTS

The author received no external funding for this work and declares no conflict of interest.

Grammarly for Education was used to correct grammatical errors and to improve a professional write-up.

CONFLICTS OF INTEREST

The authors declare that there is no conflict of interest regarding the publication of this paper

REFERENCES

- [1] U.S. Department of Energy, Grid Deployment Office. (2024). *National Transmission Planning Study*. U.S. Department of Energy, Washington, DC. Available at: <https://www.energy.gov/gdo/national-transmission-planning-study>
- [2] U.S. Energy Information Administration. (2026). *Short-Term Energy Outlook*. U.S. Energy Information Administration, Washington, DC. Available at: <https://www.eia.gov/outlooks/steo/>
- [3] North American Electric Reliability Corporation. (2025). *2025 Long-Term Reliability Assessment*. NERC, Atlanta, GA. Available at: <https://www.nerc.com/our-work/assessments/long-term-reliability-assessments>
- [4] Rand, J., Manderlink, N., Gorman, W., Wiser, R., Seel, J., Mulvaney Kemp, J., Jeong, S., & Kahrl, F. (2024). *Queued Up: 2024 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2023*. Lawrence Berkeley National Laboratory, Berkeley, CA. doi: [10.2172/2335720](https://doi.org/10.2172/2335720)
- [5] Flyvbjerg, B. (2008). Curbing optimism bias and strategic misrepresentation in planning: Reference class forecasting in practice. *European Planning Studies*, 16(1), 3-21. doi: [10.1080/09654310701747936](https://doi.org/10.1080/09654310701747936)

- [6] Flyvbjerg, B., Holm, M. K. S., & Buhl, S. L. (2002). Underestimating costs in public works projects: Error or lie? *Journal of the American Planning Association*, 68(3), 279-295. doi: [10.1080/01944360208976273](https://doi.org/10.1080/01944360208976273)
- [7] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794. doi: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785)
- [8] International Energy Agency. (2025). *Building the Future Transmission Grid*. IEA, Paris. Available at: <https://www.iea.org/reports/building-the-future-transmission-grid>
- [9] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., et al. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30, 3146-3154. Available at: https://papers.nips.cc/paper_files/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html
- [10] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774. Available at: https://papers.nips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html
- [11] Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., et al. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, 233. doi: [10.1038/s41467-019-14108-y](https://doi.org/10.1038/s41467-019-14108-y)
- [12] Poubel, R. P. B., Perpetuo, M. C. S., & Peres, W. (2025). Probabilistic transmission expansion planning under renewable and load uncertainty. *IEEE Access*, 13. doi: [10.1109/ACCESS.2025.3600866](https://doi.org/10.1109/ACCESS.2025.3600866)
- [13] Gorman, W., Mills, A., & Wiser, R. (2019). Improving estimates of transmission capital costs for utility-scale wind and solar projects to inform renewable energy policy. *Energy Policy*, 135, 110994. doi: [10.1016/j.enpol.2019.110994](https://doi.org/10.1016/j.enpol.2019.110994)
- [14] Moussa, A., Ezzeldin, M., & El-Dakhakhni, W. (2025). Machine learning and optimization strategies for infrastructure projects risk management. *Construction Management and Economics*, 43(8), 557-582. doi: [10.1080/01446193.2025.2479764](https://doi.org/10.1080/01446193.2025.2479764)
- [15] Elmousalami, H. H. (2020). Artificial intelligence and parametric construction cost estimate modeling: State-of-the-art review. *Journal of Construction Engineering and Management*, 146(1), 03119008. doi: [10.1061/\(ASCE\)CO.1943-7862.0001678](https://doi.org/10.1061/(ASCE)CO.1943-7862.0001678)
- [16] Abioye, S. O., Oyedele, L. O., Akanbi, L., Ajayi, A., Davila Delgado, J. M., et al. (2021). Artificial intelligence in the construction industry: A review of present status, opportunities and future challenges. *Journal of Building Engineering*, 44, 103299. doi: [10.1016/j.jobbe.2021.103299](https://doi.org/10.1016/j.jobbe.2021.103299)
- [17] Wood Mackenzie. (2025). *Untangling the US Transformer Supply Chain Crisis*. Wood Mackenzie, Edinburgh. Available at: <https://www.woodmac.com/press-releases/power-transformers-and-distribution-transformers-will-face-supply-deficits-of-30-and-10-in-2025/>
- [18] Lotfi, R., Sadeghi, S., Ali, S. S., Ramyar, F., Ghafourian, E., & Farbod, E. (2025). A robust, resilience machine learning with a risk approach for project scheduling. *Engineering Reports*, 7(6), e70161. doi: [10.1002/eng2.70161](https://doi.org/10.1002/eng2.70161)

- [19] Molnar, C. (2025). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable (3rd ed.)*. Independently published. Available at: <https://christophm.github.io/interpretable-ml-book/>
- [20] Wood Mackenzie. (2024). *Making the Connection: The Electric T&D Supply Chain Challenge*. Wood Mackenzie, Edinburgh. Available at: <https://www.woodmac.com/news/opinion/making-the-connection-the-electric-td-supply-chain-challenge/>
- [21] National Infrastructure Advisory Council. (2024). *Addressing the Critical Shortage of Power Transformers to Ensure Reliability of the U.S. Grid*. Cybersecurity and Infrastructure Security Agency, Washington, DC. Available at: https://www.cisa.gov/sites/default/files/2024-09/NIAC_Addressing%20the%20Critical%20Shortage%20of%20Power%20Transformers%20to%20Ensure%20Reliability%20of%20the%20U.S.%20Grid_Report_06112024_508c_pdf_0.pdf
- [22] Salih, F., & El-adaway, I. H. (2026). Integrating machine learning and reference class forecasting for construction risk contingency prediction. In *Computing in Civil Engineering 2025* (pp. 97-106). American Society of Civil Engineers, Reston, VA. doi: [10.1061/9780784486436.011](https://doi.org/10.1061/9780784486436.011)
- [23] Zhang, H., Li, R., Du, Q., Tao, J., Pineda, S., Kariniotakis, G., et al. (2025). Decision-focused learning for power system decision-making under uncertainty. *IEEE Transactions on Power Systems* (early access). doi: [10.1109/TPWRS.2025.3597806](https://doi.org/10.1109/TPWRS.2025.3597806)
- [24] Ghahramani, M., Habibi, D., & Aziz, A. (2025). A risk-averse data-driven distributionally robust optimization method for transmission power systems under uncertainty. *Energies*, 18(19), 5245. doi: [10.3390/en18195245](https://doi.org/10.3390/en18195245)
- [25] Gao, S., You, D., Chen, X., & Xie, H. (2025). Operational risk assessment of power distribution systems: A physics-informed multi-task learning approach. *International Journal of Electrical Power & Energy Systems*, 111309. doi: [10.1016/j.ijepes.2025.111309](https://doi.org/10.1016/j.ijepes.2025.111309)
- [26] Sadeghi, S. (2024). Predicting the impact of scope changes on project cost and schedule using machine learning techniques. *arXiv preprint arXiv:2412.02041*. doi: [10.48550/arXiv.2412.02041](https://doi.org/10.48550/arXiv.2412.02041)

Article Copyright: © 2026 Lukhanyiso Mpongoshé and Chidinma Queen Adieze. This is an open access article distributed under the terms of the [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use and distribution provided the original author and source are credited.

