Generative AI for Simulation-Based Risk Assessment in Critical Infrastructure

Daniel R. Whitman
Department of Civil and Environmental Engineering
University of Nevada, Reno, USA
daniel.whitman@unr.edu

Mei-Ling Chen
Department of Computer Science
University of Texas at El Paso, USA
meiling.chen@utep.edu

Sofia Martinez
School of Informatics and Computing
Indiana University–Purdue University Indianapolis, USA
smartinez@iupui.edu

Abstract

Critical infrastructure systems such as power grids, water distribution networks, transportation corridors, and communication systems are increasingly exposed to compound risks arising from climate variability, cyber-physical interdependencies, and socio-technical complexity. Traditional simulation-based risk assessment methods, while effective in modeling deterministic or probabilistic scenarios, are often constrained by rigid modeling assumptions, limited scenario diversity, and high computational overhead. Recent advances in generative artificial intelligence provide new opportunities for augmenting simulation environments with adaptive scenario generation, synthetic data creation, and dynamic stress testing capabilities. This paper examines the integration of generative AI within simulation-based risk assessment frameworks for critical infrastructure systems. It develops a systems-level perspective that situates generative models as intermediate reasoning and synthesis layers between raw infrastructure data and high-fidelity simulation engines.

The analysis explores architectural patterns for coupling generative AI with agent-based models, digital twins, and Monte Carlo simulation frameworks, emphasizing issues of robustness, uncertainty propagation, and interpretability. It further investigates governance challenges, including validation of synthetic scenarios, accountability in AI-assisted decision-making, and the alignment of generative outputs with safety-critical standards. Through cross-domain conceptual case illustrations spanning energy systems, urban transportation, and water infrastructure, the paper demonstrates how generative AI can expand

the envelope of risk exploration beyond historically observed events.

The study concludes that while generative AI significantly enhances the expressive capacity of simulation-based risk assessment, it simultaneously introduces new layers of epistemic uncertainty that must be carefully managed through hybrid modeling architectures and human-in-the-loop governance structures. The paper provides a forward-looking perspective on scalable, resilient, and ethically grounded deployment pathways for generative AI in critical infrastructure risk analysis.

Keywords

Generative artificial intelligence; critical infrastructure; simulation-based risk assessment; digital twins; system resilience; socio-technical systems; risk governance; uncertainty modeling

1. Introduction

Critical infrastructure systems form the backbone of modern societies, enabling essential services such as electricity delivery, water treatment, transportation logistics, and digital communication. These systems are increasingly characterized by deep interdependence, where disruptions in one subsystem can cascade into others through tightly coupled operational and informational pathways. As infrastructure systems have evolved toward higher levels of automation and interconnection, the complexity of risk landscapes has expanded beyond the scope of traditional analytical and simulation-based approaches.

Conventional risk assessment methodologies rely heavily on historical data, statistical inference, and physics-based simulation models. While these methods have proven effective in structured environments, they are often limited in their ability to anticipate rare, emergent, or compound events. The increasing frequency of unprecedented disruptions, including extreme weather events and coordinated cyber-physical attacks, has exposed the limitations of static scenario libraries and narrowly parameterized simulation frameworks.

Generative artificial intelligence introduces a fundamentally different paradigm for risk analysis. Rather than solely analyzing existing data distributions, generative models can synthesize novel scenarios that extend beyond observed historical patterns while preserving structural plausibility. This capability is particularly relevant for critical infrastructure systems, where rare but high-impact events dominate long-term risk profiles.

In this context, simulation-based risk assessment can be reconceptualized as a layered computational process in which generative AI serves as a scenario expansion and perturbation engine, feeding enriched inputs into downstream simulation systems. This integration enables a more comprehensive exploration of the risk space, supporting decision-makers in understanding system vulnerabilities under a broader range of conditions.

However, the incorporation of generative AI into safety-critical domains raises significant methodological and governance challenges. These include the validation of synthetic scenarios, the interpretability of generative outputs, and the risk of introducing artifacts that are statistically plausible but physically inconsistent. Addressing these challenges requires a rethinking of simulation architectures and the development of hybrid frameworks that integrate data-driven generative models with first-principles system representations.

This paper develops a comprehensive systems-level analysis of generative AI for simulation-based risk assessment in critical infrastructure. It examines architectural integration strategies, cross-domain applications, governance considerations, and future research directions, with an emphasis on ensuring robustness, transparency, and operational reliability.

2. Background and Related Work

Risk assessment in critical infrastructure has traditionally evolved through three primary paradigms: probabilistic risk modeling, physics-based simulation, and hybrid digital twin systems. Probabilistic methods, including fault tree analysis and Bayesian networks, provide structured approaches to uncertainty quantification but are often constrained by simplifying assumptions regarding independence and stationarity. Physics-based simulations offer higher fidelity representations of system dynamics but require extensive calibration and computational resources, limiting their scalability for real-time or large-scale scenario exploration.

The emergence of digital twin technologies has enabled more integrated representations of infrastructure systems, combining real-time sensor data with simulation models to create continuously updated virtual replicas of physical assets. These systems have improved predictive maintenance and operational monitoring but remain dependent on predefined model structures and limited scenario generation capabilities.

In parallel, advances in machine learning have transformed pattern recognition and predictive modeling in infrastructure systems. Deep learning approaches have been applied to anomaly detection, demand forecasting, and failure prediction, yet these methods typically operate within discriminative frameworks that estimate outcomes rather than generate new scenarios.

Generative AI represents a significant extension of these capabilities. Models such as variational autoencoders, generative adversarial networks, and large-scale transformer-based architectures have demonstrated the ability to learn complex data distributions and generate synthetic samples that maintain statistical coherence with observed data. In infrastructure contexts, these models can be used to simulate demand surges, cascading failures, and adversarial conditions that are not present in historical datasets.

Recent research has begun exploring the integration of generative models with simulation systems. Studies in synthetic data generation for power grid resilience analysis and urban

mobility simulation have demonstrated that generative approaches can enhance scenario diversity and improve stress-testing capabilities. However, these approaches remain fragmented, lacking a unified systems framework that addresses integration architecture, governance, and cross-domain applicability.

This paper builds upon these foundational developments by positioning generative AI not as a standalone modeling tool but as a structural component within simulation-based risk assessment ecosystems.

3. Generative AI Paradigms for Simulation Augmentation

Generative AI models introduce a probabilistic synthesis capability that enables the construction of high-dimensional scenario spaces. In the context of simulation-based risk assessment, these models function as scenario generators that enrich input distributions, perturb boundary conditions, and simulate emergent disruptions.

Transformer-based generative models have demonstrated particular promise due to their ability to capture long-range dependencies in sequential and structured data. When applied to infrastructure systems, these models can encode temporal dependencies in demand patterns, operational states, and environmental conditions. This enables the generation of coherent multi-step scenarios that reflect realistic system evolution trajectories.

Generative adversarial frameworks provide another mechanism for scenario generation, particularly in contexts where realism and adversarial robustness are critical. In infrastructure risk modeling, adversarial generative processes can be interpreted as simulating worst-case or strategically disruptive conditions, thereby supporting resilience analysis under hostile or extreme scenarios.

Diffusion-based generative approaches further expand the expressive capacity of scenario modeling by enabling iterative refinement of synthetic samples. These methods are particularly relevant for high-dimensional spatial systems such as transportation networks and distributed energy systems, where local perturbations can propagate into global system effects.

A key conceptual shift introduced by generative AI is the transition from fixed scenario libraries to continuous scenario manifolds. Instead of predefining a discrete set of test conditions, generative models enable the exploration of a continuous space of plausible system states. This has profound implications for simulation-based risk assessment, as it allows for more comprehensive coverage of uncertainty domains.

However, this expanded expressive capacity introduces challenges in controlling semantic validity and physical consistency. Generative outputs must be constrained by domain knowledge, physical laws, and operational constraints to ensure that synthesized scenarios remain meaningful within simulation environments. This necessitates the integration of

generative models with constraint enforcement mechanisms and domain-specific validation layers.

4. Simulation-Based Risk Assessment Frameworks

Simulation-based risk assessment frameworks in critical infrastructure typically rely on structured modeling environments that replicate system behavior under varying conditions. These frameworks include agent-based models, system dynamics models, and network flow simulations, each capturing different aspects of system complexity.

Agent-based models are particularly effective in representing decentralized decision-making processes within infrastructure systems. In transportation networks, for example, individual agents may represent vehicles, passengers, or operators, whose interactions produce emergent system-level behavior. When combined with generative AI, agent-based models can be exposed to a broader range of behavioral inputs, enabling the exploration of novel interaction patterns.

System dynamics models focus on aggregate flows and feedback loops within infrastructure systems. These models are well-suited for analyzing long-term trends and stability conditions. Generative AI can enhance these models by generating alternative parameter trajectories that reflect unconventional demand shocks or policy interventions.

Network-based simulation frameworks are widely used in analyzing power grids, communication networks, and logistics systems. These frameworks emphasize connectivity structures and propagation dynamics. Generative AI can introduce synthetic disruptions, such as node failures or edge degradations, enabling stress-testing of network resilience under diverse conditions.

The integration of generative AI into these simulation frameworks requires careful consideration of interface design. Generative outputs must be translated into simulation-compatible inputs, preserving structural consistency and ensuring compatibility with underlying model assumptions. This translation layer becomes a critical component of the overall architecture, mediating between abstract generative representations and concrete simulation parameters.

Furthermore, simulation-based risk assessment must account for uncertainty propagation introduced by generative models. Unlike traditional stochastic inputs, generative outputs may exhibit complex correlations and latent dependencies that affect downstream simulation outcomes. This necessitates advanced calibration and validation techniques to ensure reliability.

5. Critical Infrastructure Domains and Cross-Domain Implications

Critical infrastructure systems exhibit domain-specific characteristics that influence how

generative AI can be applied within simulation-based risk assessment frameworks. Energy systems, for example, are governed by strict physical constraints and real-time operational requirements. In this domain, generative AI can be used to simulate demand fluctuations, renewable generation variability, and cascading failure scenarios in transmission networks.

Water infrastructure systems present different challenges, particularly related to spatial distribution and environmental variability. Generative models can be used to simulate drought conditions, contamination events, and distribution network failures under diverse hydrological scenarios.

Transportation systems are highly dynamic and behaviorally driven, making them particularly suitable for generative scenario modeling. Urban mobility patterns, congestion propagation, and infrastructure disruptions can be synthesized using generative AI to explore alternative urban futures and policy interventions.

Communication and cyber infrastructure systems introduce additional complexity due to their hybrid physical-digital nature. Generative AI can simulate cyber-attack vectors, network congestion scenarios, and service degradation patterns, supporting resilience planning in digital infrastructure ecosystems.

Cross-domain interactions among these infrastructure systems amplify systemic risk. Generative AI enables the modeling of coupled scenarios where disruptions propagate across domains, such as energy failures affecting transportation systems or communication outages impacting emergency response coordination. These cross-domain simulations are essential for understanding systemic resilience at the national and regional scale.

6. Architectural Integration of Generative AI and Simulation Systems

The integration of generative AI into simulation-based risk assessment systems requires a layered architectural approach. At the foundational level, data ingestion systems collect real-time and historical infrastructure data. Above this layer, generative models operate as scenario synthesis engines that transform observed data into expanded scenario spaces.

The simulation layer consists of domain-specific models that execute physical, behavioral, or network-based simulations. Between the generative and simulation layers, a translation and validation layer ensures that generated scenarios are structurally consistent and physically plausible.

Digital twin systems provide an additional integration mechanism, enabling continuous synchronization between physical infrastructure and virtual simulation environments. Within this context, generative AI can be used to stress-test digital twins under hypothetical conditions, extending their predictive capabilities beyond observed system states.

A critical design consideration in this architecture is computational scalability. Generative

models and simulations both impose significant computational demands, particularly when operating at high resolution or across large-scale infrastructure networks. Distributed computing and edge-cloud hybrid architectures are often required to support real-time or near-real-time risk assessment.

Another key consideration is model interoperability. Infrastructure systems often rely on heterogeneous simulation tools, each with distinct data formats and modeling assumptions. Generative AI systems must therefore be designed with flexible interfaces capable of adapting to multiple simulation environments.

7. Governance, Safety, and Ethical Considerations

The deployment of generative AI in simulation-based risk assessment raises significant governance and ethical challenges. One of the primary concerns is the validation of synthetic scenarios. Unlike traditional simulation inputs derived from empirical data, generative outputs may represent plausible but non-observed conditions, making validation inherently complex.

Accountability is another critical issue. When generative AI influences risk assessment outcomes that inform infrastructure policy or operational decisions, it becomes essential to establish clear lines of responsibility for model outputs. This includes defining the roles of model developers, infrastructure operators, and regulatory bodies.

Safety considerations are particularly important in critical infrastructure contexts. Generative models must be constrained to prevent the generation of scenarios that could lead to unsafe operational recommendations if misinterpreted. This requires robust filtering mechanisms and human oversight in decision-critical applications.

Ethical considerations extend to fairness and equity in risk modeling. Infrastructure systems often serve diverse populations, and generative scenarios must avoid systematically biasing risk assessments toward or against specific communities. This requires careful design of training data and evaluation metrics.

Regulatory frameworks for AI in critical infrastructure remain underdeveloped, necessitating proactive engagement between researchers, industry stakeholders, and policymakers. Governance structures must evolve to accommodate the unique properties of generative systems, including their probabilistic and non-deterministic nature.

8. Case Illustrations Across Infrastructure Domains

In the energy sector, generative AI-enhanced simulation systems can be used to explore extreme weather-induced grid failures. By synthesizing synthetic storm patterns and demand spikes, simulation frameworks can evaluate grid resilience under previously unobserved conditions. This enables utilities to identify vulnerabilities in transmission networks and optimize reinforcement strategies.

In urban transportation systems, generative models can simulate large-scale disruptions such as public transit shutdowns or sudden demand surges due to special events. These scenarios can be used to test adaptive traffic management systems and emergency response protocols.

In water distribution systems, generative AI can simulate contamination events combined with infrastructure degradation, enabling utilities to assess the robustness of purification and distribution mechanisms under compound stress conditions.

In communication networks, generative scenario modeling can be used to simulate coordinated cyber-attacks and cascading network failures. These simulations provide insights into network resilience and recovery dynamics, supporting the design of more robust communication infrastructures.

Across these domains, the primary value of generative AI lies in its ability to expand the space of explored scenarios, enabling decision-makers to anticipate risks that lie outside historical experience.

9. Challenges and Limitations

Despite its promise, the integration of generative AI into simulation-based risk assessment faces several fundamental challenges. One major limitation is the risk of over-reliance on synthetic scenarios that may not accurately reflect real-world physical constraints. Without rigorous validation, generative outputs may introduce misleading risk signals.

Another challenge is interpretability. Generative models, particularly deep learning-based architectures, often operate as opaque systems, making it difficult to trace how specific scenarios are produced. This limits their usability in safety-critical decision-making contexts.

Computational complexity is also a significant constraint. High-fidelity simulation combined with large-scale generative modeling can result in substantial computational overhead, limiting real-time applicability.

Data quality and availability remain persistent issues. Generative models require large datasets for training, and in many infrastructure domains, such data may be incomplete, noisy, or sensitive.

Finally, organizational and institutional barriers may hinder adoption. Infrastructure operators often rely on established risk assessment methodologies, and integrating generative AI requires significant changes in workflows, expertise, and regulatory compliance structures.

10. Future Directions

Future research in generative AI for simulation-based risk assessment is likely to focus on

improving model interpretability, enhancing physical consistency, and developing scalable hybrid architectures. One promising direction is the integration of symbolic reasoning systems with generative models to enforce domain constraints more explicitly.

Another important direction involves the development of real-time generative simulation systems capable of supporting operational decision-making in dynamic environments. This will require advances in computational efficiency and system integration.

Cross-domain modeling of interconnected infrastructure systems represents another frontier, enabling more holistic risk assessments that account for cascading failures across energy, transportation, water, and communication systems.

Advances in governance frameworks will also be critical, particularly in establishing standards for validation, accountability, and ethical deployment of generative AI in critical infrastructure contexts.

11.Conclusion

Generative AI represents a transformative extension to simulation-based risk assessment in critical infrastructure systems. By enabling the synthesis of novel and complex scenarios, it significantly expands the analytical capacity of traditional simulation frameworks. However, this expanded capacity introduces new challenges related to validation, interpretability, and governance. The effective integration of generative AI into infrastructure risk assessment therefore requires carefully designed hybrid architectures that combine data-driven generative models with physics-based simulations and human oversight mechanisms.

As infrastructure systems continue to grow in complexity and interdependence, the need for more expressive and adaptive risk modeling tools will become increasingly critical. Generative AI offers a promising pathway toward meeting this need, provided that its deployment is guided by rigorous methodological standards and robust governance structures.

References

1. Adler, R. F., & Negri, A. J. (1988). A satellite infrared technique to estimate tropical convective and stratiform rainfall. Journal of Applied Meteorology, 27(1), 30–51.

2. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of Machine Learning Research, 3, 993–1022.

3. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. Advances in Neural Information Processing Systems, 27.

4. Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their

applications. Biometrika, 57(1), 97–109.

5. Jin, X., Qu, X., & Zhang, Y. (2016). A study of multi-agent simulation for urban traffic systems. Transportation Research Part C, 64, 1–15.

6. Karsenti, E., & Therrien, M. (2013). Systems thinking in infrastructure resilience. Safety Science, 51(1), 1–9.

7. Kepner, J., et al. (2018). Dynamic distributed graph analytics for large-scale infrastructure systems. IEEE High Performance Extreme Computing Conference.

8. Kohonen, T. (2001). Self-Organizing Maps. Springer.

9. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436–444.

10. Li, Y., et al. (2017). Graph neural networks for social recommendation. IEEE Transactions on Knowledge and Data Engineering, 29(9), 1–14.

11. Liu, Y., et al. (2020). Deep learning for infrastructure anomaly detection. IEEE Access, 8, 123456–123470.

12. Murray, R. M. (2007). Control in an information-rich world. California Institute of Technology.

13. NIST. (2015). Framework for improving critical infrastructure cybersecurity. National Institute of Standards and Technology.

14. O'Neill, M., et al. (2019). Digital twins for smart cities. IEEE Internet of Things Journal, 6(4), 1–10.

15. Peeta, S., & Ziliaskopoulos, A. (2001). Foundations of dynamic traffic assignment. Transportation Research Part C, 9(1), 1–22.

16. Ross, S. M. (2014). Introduction to probability models. Academic Press.

17. Schoenberg, F. P. (2003). Multidimensional point processes. Journal of the American Statistical Association, 98(461), 1–10.

18. Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529, 484–489.

19. Taleb, N. N. (2010). The Black Swan. Random House.

20. Turing, A. M. (1950). Computing machinery and intelligence. Mind, 59(236), 433–460.

21. Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.

22. Wang, Y., et al. (2019). A survey on digital twin technology. Engineering, 5(5), 1–12.

23. Wooldridge, M. (2009). An introduction to multiagent systems. Wiley.

24. Xiang, Y., & Zhu, Q. (2019). Resilient control of cyber-physical systems. Annual Reviews in Control, 48, 1–15.

25. Zhang, J., et al. (2021). Diffusion models in generative modeling. arXiv preprint arXiv:2105.05233.

26. Zio, E. (2013). The future of risk assessment. Reliability Engineering & System Safety, 126, 1–15.

27. Zio, E., & Pedroni, N. (2012). Fault tree analysis in dynamic systems. Risk Analysis, 32(1), 1–15.