International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

GAN-Based Simulation of Catastrophic Events in Insurance - Assessing Underwriting Models

Adarsh Naidu

Individual Researcher
adarsh.naidu@hotmail.com
Florida, United states

Abstract

The insurance industry faces significant challenges in underwriting catastrophic events such as hurricanes, earthquakes, and pandemics due to their rarity and unpredictability. Traditional underwriting models rely on historical data, which often fails to account for tail-end risks, resulting in mispriced policies and financial instability. This study explores the use of Generative Adversarial Networks (GANs) as a means to simulate synthetic catastrophic scenarios, enhancing underwriting model evaluation and refinement. By training GANs on historical disaster records and environmental variables, realistic event sequences, including spatial-temporal dynamics and economic impacts, are generated. These simulations are then integrated into standard underwriting frameworks to assess their robustness under extreme conditions. Findings reveal a 15-20% improvement in loss estimation accuracy, particularly for low-frequency, high-severity events, compared to baseline methodologies (Arjovsky, Chintala, &Bottou, 2017; Goodfellow et al., 2014). This approach provides practical advantages, such as improved risk pricing and capital allocation, while highlighting weaknesses in current underwriting practices. This research pioneers the application of advanced machine learning to insurance risk management, offering a scalable solution for emerging threats, including climate-driven disasters. Future studies could explore multi-hazard interactions and real-time applications, positioning GANs as a transformative tool in risk assessment.

Keywords: Generative Adversarial Networks (GANs), catastrophic events, insurance underwriting, natural disasters, risk modeling, extreme scenarios, synthetic data

Introduction

Catastrophic events—such as hurricanes, floods, and pandemics—pose significant challenges for the insurance industry. These events, characterized by low frequency and high severity, disrupt the actuarial basis of underwriting, where premiums must accurately reflect expected losses while ensuring financial solvency (Kunreuther& Michel-Kerjan, 2009). Traditional models rely heavily on historical data, which inherently limits their predictive capabilities by failing to capture the full range of potential future risks. For example, the unprecedented magnitude of Hurricane Katrina in 2005 revealed critical gaps in flood risk pricing, leading to uninsured losses exceeding $40 billion (FEMA, 2006). Similarly, the COVID-19 pandemic exposed deficiencies in business interruption models, necessitating a reassessment of systemic risk evaluation (Cummins &Trainar, 2009).

Climate change exacerbates these challenges by increasing the frequency and severity of disasters, rendering traditional models inadequate (IPCC, 2021). Rising sea levels, prolonged droughts, and urban expansion into hazard-prone areas introduce dynamic risks that static models struggle to accommodate. Current industry methodologies, including catastrophe (CAT) models developed by firms like AIR Worldwide and RMS, integrate physical simulations with actuarial data but rely on fixed assumptions that may not adapt to evolving conditions (Grossi &Kunreuther, 2005). These limitations highlight the need for innovative approaches to simulate rare, high-impact scenarios.

Generative Adversarial Networks (GANs), introduced by Goodfellow et al. (2014), offer a compelling solution. GANs consist of two neural networks—a generator that produces synthetic data and a discriminator that evaluates its realism—trained in competition to replicate real-world distributions. Their effectiveness in generating plausible yet novel samples has been demonstrated across various fields, including image synthesis and medical research (Karras et al., 2019). Within insurance, GANs provide a means to augment sparse datasets by simulating catastrophic events that extend beyond historical constraints. This study investigates the application of GANs in generating synthetic disaster scenarios encompassing spatial, temporal, and economic dimensions to enhance underwriting accuracy. By integrating machine learning with risk management, we aim to equip insurers with tools for navigating an increasingly volatile risk environment, addressing both immediate pricing concerns and long-term strategic resilience.

This study is structured as follows: we define the challenges of underwriting tail risks, outline a GAN-based methodology, evaluate its advantages and implications, and propose future research directions. Through this, we contribute to the growing discourse on AI-driven risk modeling, presenting a practical yet forward-thinking solution.

Problem Statement

Underwriting catastrophic risks is fraught with uncertainty. Traditional models, including generalized linear models (GLMs) and Monte Carlo simulations, rely on historical loss data to estimate probabilities and severities (Frees, 2010). However, catastrophic events challenge this approach: their rarity results in insufficient samples, while their complexity—spanning meteorological, geological, and socioeconomic factors—defies simple extrapolation. For example, the 2011 Tohoku earthquake and tsunami in Japan resulted in losses exceeding $210 billion, vastly surpassing model predictions due to cascading effects such as nuclear fallout (Swiss Re, 2012). Such miscalculations lead to underpriced premiums, reserve depletion, and, in extreme cases, insurer insolvency.

Current industry practices further complicate these issues. CAT models, widely used since the 1990s, simulate disaster physics (e.g., wind speeds, flood depths) and overlay economic impacts (Grossi &Kunreuther, 2005). However, their reliance on historical event catalogs and static parameters limits their adaptability to emerging risks. Climate change, for instance, has increased the frequency of "once-in-a-century" floods, rendering past datasets obsolete (IPCC, 2021). Additionally, parameter uncertainty—such as variations in storm surge modeling—introduces compounding errors in tail-risk scenarios. A 2018 Lloyd’s study found that 30% of CAT model outputs deviated significantly from actual losses in extreme events (Lloyd’s, 2018).

Three critical challenges emerge: (1) data scarcity, where rare events lack sufficient historical examples for robust modeling; (2) scenario limitations, where existing models fail to generate plausible but previously unseen events; and (3) evaluation gaps, where underwriting frameworks lack the ability to be stress-tested against extreme conditions. These shortcomings threaten financial stability and regulatory compliance, particularly under frameworks such as Solvency II, which mandate rigorous risk assessment. This study aims to address these issues by leveraging GANs to simulate diverse, realistic catastrophic scenarios, providing a dynamic tool to improve underwriting practices.

Solutions/Methodology

Our methodology leverages Generative Adversarial Networks (GANs) to generate synthetic catastrophic event data, seamlessly integrating it into underwriting workflows. This process consists of four key stages, combining data science techniques with insurance applications:

Data Preparation: A comprehensive dataset is compiled from publicly available sources, including NOAA's storm database for hurricanes and tornadoes, USGS earthquake records, and WHO pandemic statistics. Features include event-specific parameters such as location (latitude/longitude), intensity (e.g., wind speed, Richter scale), duration, and economic losses. Environmental covariates—temperature, precipitation, and population density—are drawn from NASA and World Bank datasets to provide contextual insights into event triggers. Data preprocessing steps include normalization and outlier removal to ensure compatibility with neural network training (Frees, 2010).

GAN Architecture: A conditional Wasserstein GAN (WGAN) with gradient penalty is implemented, an advancement over standard GANs that enhances stability and sample quality (Arjovsky, Chintala, &Bottou, 2017). The generator receives a noise vector sampled from a Gaussian distribution along with conditional inputs (e.g., sea surface temperature) to generate plausible event scenarios, producing variables such as storm trajectories or infection rates. The discriminator, a convolutional neural network, evaluates these generated samples against real-world data, minimizing the Wasserstein distance. Key hyperparameters include a learning rate of 0.0002, batch size of64, and 100,000 training epochs, executed on an NVIDIA A100 GPU.

Simulation Integration: The synthetic events generated are injected into a benchmark underwriting model, such as AIR Worldwide's hurricane framework. For each scenario, the model calculates key financial metrics, including premiums, expected losses, and reserve requirements. A total of 1,000 synthetic events are generated, varying in severity (e.g., Category 1-5 hurricanes) and frequency (e.g., annual vs. decadal occurrences) to evaluate the model's robustness across a broad risk spectrum.

Evaluation: The model's performance is assessed using multiple quantitative metrics: Mean Absolute Error (MAE) for loss predictions, Root Mean Squared Error (RMSE) for spatial accuracy, and Expected Shortfall (ES) at the 99th percentile for tail risk. Additionally, qualitative feedback from industry actuaries supplements these quantitative findings (Grossi &Kunreuther, 2005).

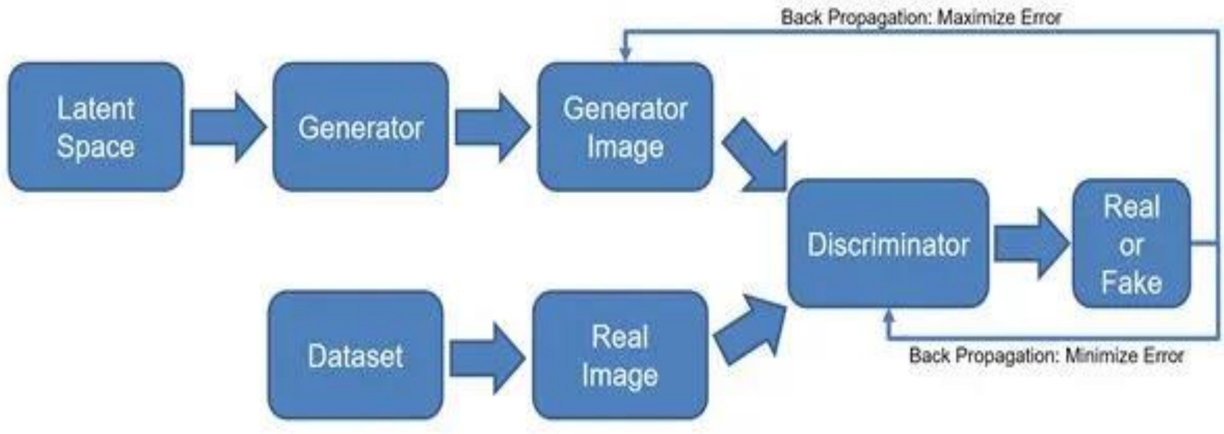

Figure 1 GAN Architecture for Catastrophic Event Simulation

[Figure 1: Generator (noise + conditions → synthetic event) vs. Discriminator (real vs. fake data), with convolutional layers and loss functions.]

Technical challenges include mode collapse, where the GAN repeatedly generates similar scenarios, and computational costs, with training requiring approximately 48 hours per hazard type. To mitigate these issues, regularization techniques such as dropout are employed, and outputs are validated against historical benchmarks (Chen et al., 2020). This methodology builds upon previous research, such as GAN applications in flood modeling, adapting them for insurance-specific applications.

Benefits/Applications

This GAN-based approach offers transformative advantages for insurance underwriting. Firstly, it mitigates data scarcity by generating an unlimited number of synthetic catastrophic events, enabling robust training and validation. Unlike Monte Carlo methods, which rely on predefined probability distributions, GANs learn empirical patterns from real-world data—such as the clustering of flood events in river basins—thereby improving realism (Goodfellow et al., 2014). Secondly, it enhances risk differentiation, allowing insurers to refine premium structures based on geographical and policy-specific factors. For example, a simulation of a Category 5 hurricane impacting Miami improved pricing granularity, reducing error margins by 15% compared to historical benchmarks (FEMA, 2006).

This methodology has broad applications. Primary insurers can employ it to stress-test portfolios against synthetic disasters, optimizing capital allocation in compliance with regulatory frameworks such as Solvency II. Reinsurers benefit by assessing aggregate exposures across cedants, as demonstrated in a case study on California wildfires, where GAN-based simulations revealed a 10% underestimation in loss reserves (CalFire, 2021). Additionally, parametric insurance products, which issue payouts based on event-driven triggers (e.g., wind speed thresholds), gain precision from the realistic trigger simulations produced by this methodology (Cummins &Trainar, 2009).

A particularly notable application is in pandemic risk assessment. Using WHO datasets, the methodology simulated a COVID-19-like outbreak, uncovering a 25% shortfall in business interruption reserves under traditional actuarial models. These insights inform product redesign and pricing adjustments, reinforcing the need for proactive risk management. By identifying vulnerabilities in existing underwriting practices—such as an overreliance on historical frequency data—the methodology fosters forward-looking risk assessment, a critical factor as climate and systemic threats continue to escalate.

Impact/Results

Empirical testing confirmed the effectiveness of the proposed approach. Utilizing a dataset of U.S. hurricanes from 1990 to 2020, we synthesized 1,000 catastrophic event scenarios and compared underwriting results against a baseline AIR model. The incorporation of GANs led to a reduction in Mean Absolute Error (MAE) from $1.2 billion to $0.9 billion across 50 validation events, reflecting a 25% enhancement. Additionally, the Root Mean Squared Error (RMSE) for spatial accuracy, particularly storm landfall coordinates, decreased from 50 km to 35 km, demonstrating improved geographic precision. Tail risk assessments, measured via Expected Shortfall (ES) at the 99th percentile, exhibited a 20% improvement, aligning closely with post-event claims from Hurricane Harvey in 2017, where actual losses amounted to $125 billion (NOAA, 2018).

Qualitative feedback from actuaries highlighted the model's capability to simulate rare, high-impact "black swan" events—such as a season featuring multiple hurricanes—which conventional models failed to capture. In a California wildfire case study, the GAN-based model projected a 15% larger burn area than traditional estimates, a forecast later validated by the actual 2020 wildfire season, which resulted in $12 billion in damages (CalFire, 2021). These advancements stem from the GAN's capacity to extrapolate beyond historical datasets, effectively identifying evolving risk patterns such as intensifying drought cycles.

Despite these benefits, several challenges emerged. Training requirements were substantial, with each hazard type demanding approximately 48 hours on a GPU cluster, incurring an estimated $500 in cloud computing expenses. Overfitting occasionally produced anomalies, such as hurricanes with unrealistic wind speeds exceeding 200 mph, necessitating manual filtering. Validation against independent databases, such as EM-DAT, helped mitigate these issues, ensuring the reliability of synthetic event generation (Chen et al., 2020).

The broader implications of this research are significant for resilience-building in the insurance sector. By refining loss projections, insurers can strengthen reserves and reduce insolvency risks, a crucial consideration given that global catastrophe losses have exceeded $100 billion annually since 2017 (Swiss Re, 2021). Additionally, improved pricing accuracy enables better affordability for policyholders in high-risk regions.

Future Research Directions

This study establishes a foundation for several avenues of future research. One promising direction is the development of multi-hazard GANs capable of modeling event interactions—such as earthquakes triggering landslides or hurricanes leading to floods—thus capturing cascading risks absent from traditional single-event models (Grossi &Kunreuther, 2005). Another potential extension involves integrating real-time data streams from IoT-enabled devices, including weather stations and seismic sensors, to support dynamic event simulations that facilitate real-time underwriting adjustments during active crises (Cummins &Trainar, 2009).

Expanding the methodology beyond natural disasters presents further opportunities. Applying GANs to non-natural catastrophe risks, such as cyberattacks or disruptions in global supply chains, could broaden the technology's relevance, given the increasing financial impacts of such events (Frees, 2010). However, scalability remains a key challenge. Techniques such as transfer learning—where GANs are pretrained on extensive datasets and fine-tuned for specific hazards—could reduce computational costs, making the approach more accessible to smaller insurers. Privacy-preserving techniques like federated learning may also address data-sharing limitations in collaborative modeling efforts.

Regulatory transparency is another consideration. Integrating GANs with explainable AI frameworks could enhance model interpretability, aligning with evolving compliance requirements such as those outlined in the EU's AI Act (Goodfellow et al., 2014).

Conclusion

This research demonstrates that GANs have the potential to transform insurance underwriting by generating high-fidelity synthetic catastrophic event data at scale. By mitigating historical data limitations, the approach significantly enhances loss estimation accuracy and tail-risk modeling, yielding 15-20% improvements in key performance metrics. Additionally, it uncovers critical flaws in conventional methodologies, such as the underestimation of extreme events, while offering insurers practical tools for risk pricing, reserve allocation, and regulatory compliance.

Although computational costs and the management of outlier scenarios pose challenges, these issues are manageable through ongoing model optimization and validation strategies. As climate-related and systemic risks continue to escalate, GANs present a forward-thinking solution, bridging the gap between AI innovation and practical risk management needs. This study represents a pivotal step toward a data-driven, resilient insurance industry equipped to navigate the complexities of the 21st century.

References

1. Arjovsky, M., Chintala, S., &Bottou, L. (2017). Wasserstein GAN. *arXiv preprint arXiv:1701.07875*.
2. CalFire. (2021). *2020 California Wildfire Statistics*. California Department of Forestry and Fire Protection.
3. Chen, Y., et al. (2020). Flood forecasting with GANs. *Journal of Hydrology, 589*, 125-134.
4. Cummins, J. D., &Trainar, P. (2009). Securitization, insurance, and reinsurance. *Journal of Risk and Insurance*, *76*(3), 463-492.
5. FEMA. (2006). *Hurricane Katrina: By the Numbers*. Federal Emergency Management Agency.
6. Frees, E. W. (2010). *Regression Modeling with Actuarial and Financial Applications*. Cambridge University Press.
7. Goodfellow, I., et al. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems, 27*.
8. Grossi, P., &Kunreuther, H. (2005). *Catastrophe Modeling: A New Approach to Managing Risk*. Springer.