

Generative AI In Finance And Banking

Interview Preparation Guide



Visit Our Website
www.gsdccouncil.org



1. What is Generative AI, and how is it different from predictive AI?

Generative AI learns patterns from historical data to create new, realistic outputs, while predictive AI focuses on forecasting future outcomes. In finance, generative AI can produce synthetic stock price sequences, rare event simulations, or customer profiles, which helps with stress testing, scenario analysis, and model robustness. Predictive AI, on the other hand, would only predict the next-day stock price or credit default based on historical trends.

2. Explain the concept of GANs. What are their components?

A GAN (Generative Adversarial Network) has two neural networks: the Generator, which creates synthetic data, and the Discriminator, which evaluates if the data is real or fake. They train in competition, improving each other iteratively. In finance, GANs can simulate market scenarios, generate synthetic fraud data, or create alternate historical price paths to test portfolio strategies.

3. What is a Variational Autoencoder (VAE)? How does it differ from a GAN?

A VAE encodes data into a latent probabilistic space and reconstructs it via decoding. Unlike GANs, which prioritize realistic output through adversarial training, VAEs focus on structured latent representations, making them suitable for anomaly detection, scenario simulations, and explainable modeling in finance.

4. Why is explainability important in financial AI models?

Explainability ensures regulatory compliance, trust, and transparency. Financial institutions must justify decisions in credit scoring, risk modeling, and fraud detection. Tools like SHAP and LIME help interpret AI outputs by showing feature contributions and model reasoning.

5. What are the key metrics in credit risk scoring?

Answer: Key metrics include:

- PD (Probability of Default): Likelihood a borrower will default.
- LGD (Loss Given Default): Potential loss in case of default.
- EAD (Exposure at Default): Exposure amount at the time of default.

6. How is NLP applied in finance?

Answer: NLP processes unstructured text such as earnings call transcripts, news articles, or analyst reports. Applications include sentiment analysis to forecast market reactions, extracting key insights for investment decisions, and detecting potential risks. For example, negative sentiment in an earnings call may signal a likely drop in stock price.

7. Define Agentic AI and its potential applications in banking.

Answer: Agentic AI refers to systems that act autonomously to achieve goals. In banking, this could include autonomous trading bots, portfolio rebalancing agents, and automated risk monitoring systems. It combines generative AI with reinforcement learning to make decisions without human intervention.

8. If historical market data is limited, how would you simulate future stock prices?

Answer: I would use GANs or TimeGANs to generate synthetic price sequences that mimic historical statistical properties. These sequences can be used for stress testing strategies, scenario analysis, and backtesting portfolios. Validation involves comparing distributions with real data to ensure accuracy.

9. How would you detect fraudulent transactions using generative AI?

Answer: I would train an autoencoder or GAN on normal transactions. High reconstruction error (autoencoder) or discriminator flags (GAN) indicate anomalies. Metrics like Precision, Recall, and F1-score help assess performance. GANs can also generate synthetic fraud cases to improve training with rare events.

11. How would you forecast stock trends using textual news data?

I would extract sentiment scores from news articles and social media using NLP models like FinBERT. These scores are combined with historical price data and fed into LSTM or attention-based neural networks to forecast trends. This approach captures market sentiment along with quantitative trends.

12. How would you ensure an AI trading bot complies with regulations while acting autonomously?

Implement risk constraints and supervisory layers, maintain decision logs for auditability, and ensure regular validation and scenario testing. Compliance requires integrating rules to prevent excessive risk-taking while still allowing autonomous execution.

13. Which tools and libraries are used for generative AI in finance?

- Deep Learning: PyTorch, TensorFlow, Keras
- NLP: Hugging Face Transformers, SpaCy, FinBERT
- Backtesting/Trading: QuantConnect, Zipline
- Machine Learning: Scikit-learn, XGBoost
- Explainability: SHAP, LIME
- Time Series: Prophet, TimeGAN

14. How do you preprocess financial time series data for LSTM?

Normalize or standardize the data, handle missing values, create sequential windows of observations, split data into training/validation/test sets, and optionally generate technical indicators like moving averages or volatility.

15. What challenges exist when training GANs on financial data?

Challenges include mode collapse, unstable training, overfitting due to small datasets, and non-stationarity of financial data. Solutions involve proper loss function tuning, regularization, and combining synthetic with real data.

16. Explain a project where you implemented generative AI in finance.

I used a GAN to generate synthetic stock prices for backtesting a trading strategy under rare conditions. The generated data was validated statistically against historical data. This approach helped evaluate portfolio risk in scenarios that hadn't occurred historically.

17. How do you evaluate synthetic data generated by a GAN?

By statistical similarity to real data, downstream model performance, and assessing diversity/realism. For finance, testing synthetic sequences in backtested trading strategies or risk models ensures they are actionable.

18. If a fraud detection model has high false positives, what steps would you take?

High false positives indicate that the model is flagging too many legitimate transactions as fraudulent, which can annoy customers and waste operational resources. First, I would analyze the feature set using explainability tools like SHAP or LIME to see which features contribute to false positives. Then, I would rebalance the dataset using techniques like SMOTE or GAN-based augmentation to improve rare fraud detection while reducing overfitting. Adjusting the decision threshold and testing with a validation set ensures an optimal trade-off between precision and recall. Finally, I would consider ensemble models or rule-based hybrid systems to further reduce false alarms.

19. How can generative AI be used to model market stress scenarios?

Stress scenarios are rare but impactful events that historical data may not cover. Generative AI models like TimeGAN or VAE can simulate synthetic sequences representing extreme market fluctuations, including high volatility or sudden crashes. These synthetic scenarios can be used to stress-test portfolios, risk models, and trading strategies, allowing risk managers to anticipate potential losses. For instance, a bank could simulate a 2008-style market crash and evaluate Value at Risk (VaR) or Expected Shortfall (CVaR) for its investment portfolio.

20. What are the risks of using generative AI in finance?

Risks include generating unrealistic or biased data, overfitting to synthetic patterns, and regulatory non-compliance. Misinterpreted synthetic outputs could lead to wrong investment or credit decisions. To mitigate these risks, it's essential to validate generated data against historical distributions, use explainable AI tools for interpretability, and maintain human oversight in decision loops. Additionally, documentation of model assumptions and stress-testing synthetic outputs can ensure regulatory transparency.

21. Where do you see generative AI in finance in the next 5 years?

Generative AI will likely become integral to risk management, scenario simulation, fraud detection, and autonomous trading. Agentic AI systems may independently rebalance portfolios based on evolving market conditions. NLP-powered generative AI could analyze earnings calls, regulatory filings, and news in real time to guide investment decisions. Synthetic data generation will help backtest strategies for rare events, improving regulatory compliance and financial resilience.

22. How might agentic AI change the role of financial analysts?

Agentic AI can automate data analysis, scenario generation, and even initial investment recommendations, allowing analysts to focus on strategic oversight, validation, and high-level decision-making. For example, in credit risk assessment, an agentic system could automatically generate borrower profiles and score them, while analysts interpret results, manage exceptions, and make judgment calls. Analysts become decision supervisors rather than data processors, increasing efficiency and reducing human error.

23. How would you combine generative AI with reinforcement learning in finance?

Generative AI can simulate multiple market conditions, including rare or extreme events, to train reinforcement learning (RL) agents. The RL agent learns optimal trading or portfolio strategies by interacting with synthetic environments, exploring actions without real-world risk. For example, a trading bot could test high-volatility market responses generated by a GAN, optimizing portfolio allocation while learning risk limits. This combination allows robust, autonomous decision-making in dynamic financial markets.

24. Explain GAN and VAE loss functions.

GAN Loss: The generator tries to minimize the discriminator's ability to distinguish real vs. fake: $\min_G \max_D V(D,G) = E[\log D(x)] + E[\log(1-D(G(z)))]$. In finance, tuning this loss ensures synthetic data closely mimics historical market data.

VAE Loss: Combines Reconstruction Loss (how accurately the input is reproduced) and KL Divergence (ensuring latent space follows a normal distribution). VAEs are used in anomaly detection and scenario generation, where preserving underlying statistical structures is crucial.

25. What metrics are used for regression models in finance?

Metrics include RMSE (Root Mean Squared Error), MAE (Mean Absolute Error), and MAPE (Mean Absolute Percentage Error). For example, when predicting daily stock prices, RMSE measures deviation magnitude, MAE measures average prediction error, and MAPE expresses accuracy in percentage terms. These metrics help tune model performance and compare different models.

26. What metrics are used for classification problems like fraud detection?

Important metrics include Accuracy, Precision, Recall, F1-score, and AUC-ROC. Since fraud is rare, Precision (avoiding false alarms) and Recall (detecting actual fraud) are particularly critical. AUC-ROC evaluates overall discrimination ability. For example, in credit card fraud detection, prioritizing high recall ensures fraud cases are not missed, while precision reduces customer inconvenience.

27. How do you validate a credit scoring model?

Use out-of-sample testing, cross-validation, and stress-testing under extreme scenarios. Evaluate PD, LGD, and EAD metrics for accuracy. Explainable AI tools like SHAP can show feature contributions, ensuring regulatory transparency. Sensitivity analysis helps understand how changes in borrower behavior affect predicted risk, critical for Basel compliance.

28. Your synthetic dataset shows unrealistic market spikes. How do you fix it?

Answer: Unrealistic spikes may result from GAN instability or noise in input data. Solutions include adjusting learning rates, refining architecture, adding regularization, or blending synthetic with historical data. Validation through statistical measures like Kolmogorov-Smirnov tests ensures generated sequences match real market distributions.

29. You notice a model predicts high PD for low-risk borrowers. What steps would you take?

Answer: Investigate feature importance using SHAP/LIME, check for data preprocessing errors, retrain with a balanced dataset, and perform cross-validation. Additionally, verify that macro-economic variables are properly included to prevent biased predictions.

30. How would you detect rare fraudulent transactions using limited labeled data?

Answer: Use GANs to generate synthetic fraud samples, increasing the diversity of the minority class. Combine with anomaly detection models like autoencoders. Evaluate using Precision-Recall curves, focusing on capturing rare events while minimizing false positives.

31. How can generative AI help in scenario planning for portfolio risk?

Answer: Generate synthetic market sequences, including extreme volatility or rare events. Test portfolios under these scenarios to measure VaR, CVaR, and stress resilience. For example, simulate a market crash to check whether derivatives hedges perform as expected and adjust portfolio allocation proactively.

32. Describe a time you used AI to solve a financial problem.

Answer: I applied NLP to earnings call transcripts to extract sentiment signals, combined them with historical stock prices, and trained an LSTM to predict short-term trends. This approach improved trading strategy accuracy and provided actionable insights that weren't captured by quantitative data alone.

33. How do you stay updated with advances in Generative AI for finance?

I stay updated through several methods:

- Following arXiv papers
- Participating in Kaggle competitions
- Attending industry webinars
- Reading AI-finance blogs
- Experimenting with models on real datasets
- Completing certifications and online courses to understand regulatory impacts, new architectures, and best practices for financial applications.

34. How would you explain Generative AI to a non-technical banking manager?

"Generative AI learns patterns from historical data to create realistic synthetic datasets or scenarios. In banking, it helps us simulate market conditions, predict risk, detect fraud, and test strategies without relying solely on past events, allowing safer and faster decision-making."

35. What are some ethical considerations when using Generative AI in finance?

Ethical considerations include:

- Avoiding bias in synthetic data
- Ensuring transparency and fairness
- Maintaining regulatory compliance
- Preventing misuse of sensitive information
- Implementing human oversight, monitoring, and explainability to ensure ethical and responsible deployment.

36. How will Agentic AI and Generative AI reshape financial services?

Automation of trading, portfolio management, fraud detection, and risk assessment will become standard. Analysts will focus on strategy, validation, and oversight, while AI handles repetitive tasks. Decision-making becomes faster, data-driven, and more resilient, especially under rare market events.

37. How can banks integrate Generative AI with traditional ML systems?

Generative models can augment datasets or create stress scenarios, which are then fed into traditional ML models to improve predictions and risk assessment. For instance, synthetic loan profiles could improve credit scoring model robustness or help backtest portfolio performance under extreme market conditions.

38. What is the biggest challenge in implementing Generative AI in banking?

The main challenges are model reliability, explainability, and regulatory compliance, especially with volatile, non-stationary financial data. Mitigation includes extensive validation, explainable AI, human supervision, and documentation for audit purposes.

39. How can NLP insights complement numerical financial models?

NLP provides leading indicators from news, social media, and reports that numerical data may miss. For example, sentiment from analyst reports can enhance stock price prediction models, helping portfolio managers anticipate market reactions before they are reflected in prices.

40. If you had to choose one area to start applying Generative AI in finance, which would it be and why?

I would start with fraud detection or scenario planning. Limited historical data for rare events makes synthetic data generation highly valuable. It improves risk mitigation, model robustness, and regulatory compliance, while also providing practical insights for portfolio and operational risk management.

CERTIFICATION IN GENERATIVE AI IN FINANCE AND BANKING



ABOUT GSDC CERTIFICATION



EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.



CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now