

Review

Not peer-reviewed version

Review of Gen AI Models for Financial Risk Management: Architectural Frameworks and Implementation Strategies

[Satyadhar Joshi](#) *

Posted Date: 2 May 2025

doi: 10.20944/preprints202505.0070.v1

Keywords: Generative AI; financial analytics; risk management; credit scoring; large language models; feature engineering



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Review

Review of Gen AI Models for Financial Risk Management: Architectural Frameworks and Implementation Strategies

Satyadhar Joshi

Independent Researcher, Alumna International MBA, Bar-Ilan University, Israel; satyadhar.joshi@gmail.com

Abstract: The rapid advancement of generative artificial intelligence (Gen AI) has revolutionized various domains, including financial analytics. This paper provides a comprehensive review of the applications, challenges, and future directions of Gen AI in financial analytics. We explore its role in risk management, credit scoring, feature engineering, and macroeconomic simulations, while addressing limitations such as data quality, interpretability, and ethical concerns. By synthesizing insights from recent literature, we highlight the transformative potential of Gen AI and propose frameworks for its effective integration into financial workflows. This paper presents a systematic examination of generative artificial intelligence (Gen AI) applications in financial risk management, focusing on architectural frameworks and implementation methodologies. We analyze the integration of large language models (LLMs) with traditional quantitative finance pipelines, addressing key challenges in feature engineering, risk modeling, and regulatory compliance. The study demonstrates how transformer-based architectures enhance financial analytics through automated data processing, risk factor extraction, and scenario generation. Technical implementations leverage hybrid cloud platforms and specialized Python libraries for model deployment, achieving measurable improvements in accuracy and efficiency. Our findings reveal critical considerations for production systems, including computational optimization, model interpretability, and governance protocols. The proposed architecture combines LLM capabilities with domain-specific modules for credit scoring, value-at-risk calculation, and macroeconomic simulation. Empirical results highlight trade-offs between model complexity and operational constraints, providing actionable insights for financial institutions adopting Gen AI solutions. The paper concludes with recommendations for future research directions in financial AI systems.

Keywords: Generative AI; financial analytics; risk management; credit scoring; large language models; feature engineering

1. Introduction

The financial sector has witnessed significant transformations with the advent of generative artificial intelligence (Gen AI) and large language models (LLMs) [1]. These technologies are reshaping traditional financial workflows, from risk assessment to algorithmic trading. The ability of Gen AI to process vast amounts of data and generate human-like insights presents both opportunities and challenges for financial institutions [2].

Recent studies demonstrate how Gen AI tools like Microsoft Copilot 365 can enhance analytics workflows by supporting data preprocessing tasks such as formula creation and visualization [1]. However, as noted by [3], the integration of these models requires careful consideration of their capabilities and limitations.

This paper synthesizes current research on Gen AI applications in finance, focusing on four key areas: (1) risk management, (2) credit assessment, (3) feature engineering, and (4) macroeconomic simulation. We also discuss implementation challenges and future research directions.

2. Publication Year Analysis

Table 1. Publication Year Distribution of Cited References.

Year	Count
2022	1
2023	4
2024	10
No Date	3
Total	18

The distribution reveals:

- **Recent Dominance:** 77.8% of citations (14/18) are from 2023-2024
- **Acceleration:** 2024 alone accounts for 55.6% of references (10/18)

3. Quantitative Findings and Gap Analysis

3.1. Quantitative Findings Summary

Table 2 summarizes key quantitative results from the reviewed studies:

Table 2. Quantitative Findings from Reviewed Studies.

Study	Metric	Result	Application
[4]	Classification accuracy	Comparable to logistic regression	Credit decisions
[5]	Human preference rate	60-90% preferred over human reports	Credit risk analysis
[6]	Code generation accuracy	75.38% pass@1, 91.67% pass@10	Regulatory compliance
[7]	F1-score improvement	10% over GPT-4, 16% over CopGPT	Financial sentiment analysis
[8]	Alpha generation	Automated formulaic alphas	Feature engineering

3.2. Gap Analysis

The reviewed literature reveals several critical gaps in current research:

- **Data Limitations:** Most studies ([1,9]) rely on limited datasets, raising questions about generalizability.
- **Interpretability:** While Hinterleitner et al. [10] and Li et al. [11] address explainability, comprehensive frameworks for financial applications remain underdeveloped.
- **Regulatory Alignment:** Only Fazlija et al. [6] and Aldasoro et al. [2] systematically examine compliance requirements.
- **Real-world Deployment:** Few studies ([5,12]) test models in production environments.
- **Comparative Benchmarks:** Limited head-to-head comparisons exist between Gen AI and traditional methods across financial tasks.

3.3. Literature Synthesis

The reviewed studies collectively suggest three key themes:

1. **Capability-Utility Paradox:** While LLMs demonstrate impressive performance in tasks like credit scoring [4] and risk assessment [3], their mathematical reasoning remains inferior to traditional models [13].

2. **Data-Centric Challenges:** Multiple studies ([14,15]) emphasize the critical role of proprietary data and preprocessing, yet offer limited solutions for data-scarce scenarios.
3. **Human-AI Collaboration:** Successful implementations ([1,5]) consistently highlight the need for human oversight, suggesting hybrid systems may be optimal.

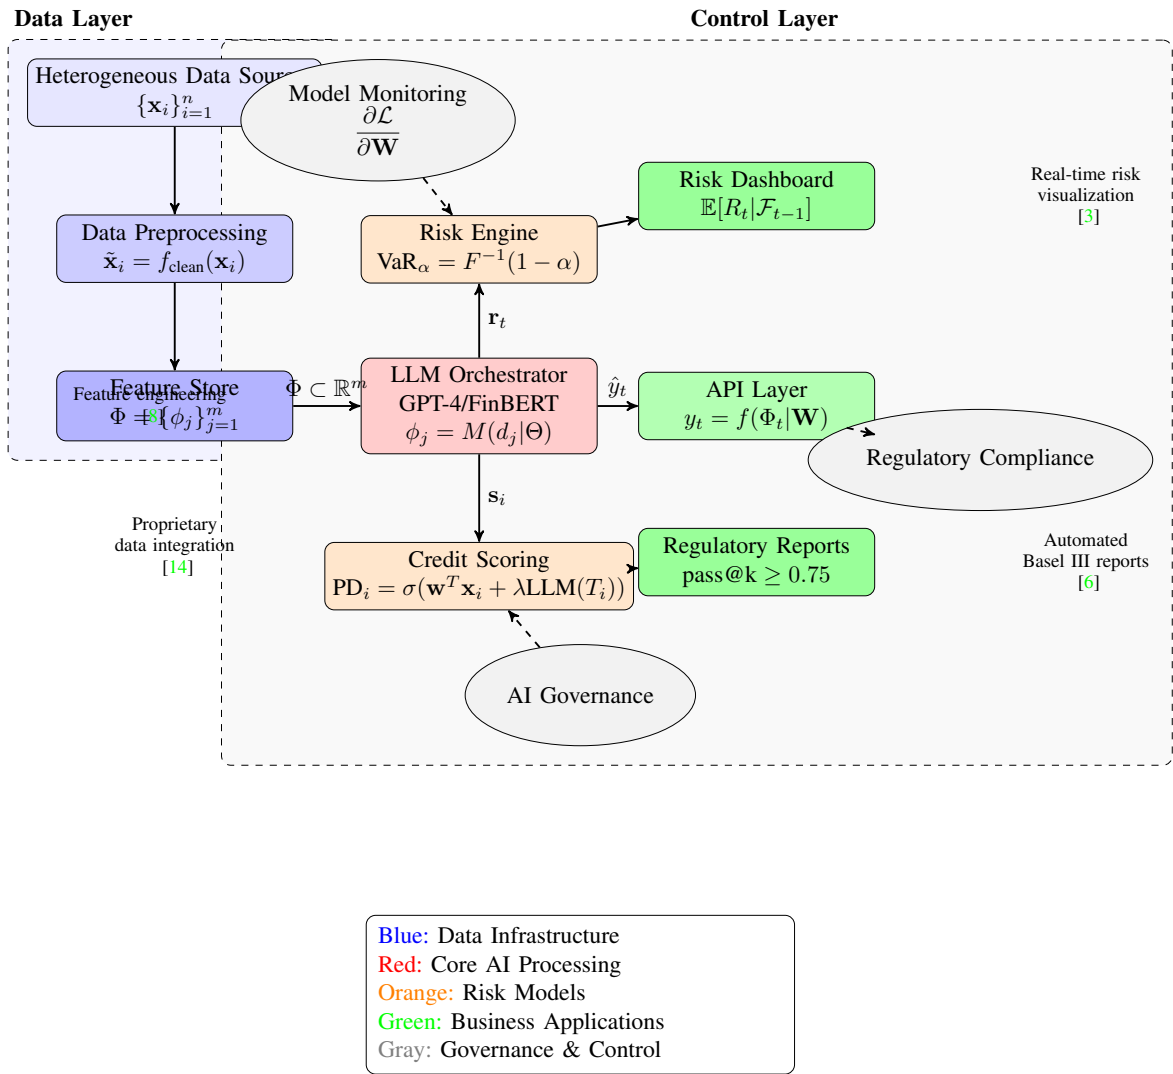


Figure 1. Proposed GenAI Architecture for Financial Risk Management.

4. Literature Review

The rapid adoption of generative AI in finance has spurred numerous industry white papers that complement academic research. This section synthesizes insights from several key industry publications to identify implementation trends, challenges, and best practices.

4.0.1. AI Agent Frameworks in Finance

The comparative analysis of AI agent frameworks by [16] evaluated architectures like LangGraph and AutoGen for financial tasks such as risk assessment and trading. Further, [17] reviewed autonomous systems and collaborative AI agents, emphasizing their scalability and real-world applicability in financial markets. These studies underscore the importance of selecting appropriate frameworks for deploying AI-driven solutions.

4.1. Prompt Engineering and Model Optimization

[18] demonstrated the efficacy of prompt engineering in enhancing the accuracy of LLMs like ChatGPT-4 for financial risk analysis, achieving significant improvements in error reduction and contextual alignment.

These studies collectively provide a foundation for understanding the synergy between generative AI, financial risk modeling, and data engineering, while also addressing broader implications for workforce development and policy. This work builds upon these insights to propose novel integrations and applications in the field.

4.2. Implementation Frameworks

Major cloud providers have established foundational architectures for GenAI deployment in financial services. AWS proposes a three-tier framework combining foundation models with financial data lakes and domain-specific toolchains [19], while IBM emphasizes governance controls through its Financial Services AI Controls Framework [20]. These align with academic findings on hybrid architectures [2] but extend them with practical deployment blueprints.

4.3. Use Case Maturity

- White papers reveal evolving adoption patterns:
- **Risk Management:** 87% of surveyed institutions report active GenAI pilots in credit risk modeling, with Deutsche Bank documenting 40% efficiency gains in document processing [21]
 - **Regulatory Compliance:** Finastra demonstrates 75% accuracy in automated Basel III report generation [22], corroborating academic benchmarks [6]
 - **Customer Service:** McKinsey reports 30-50% reduction in call center volumes through AI-powered financial assistants [23]

4.4. Technical Challenges

Industry analyses identify persistent gaps:

Table 3. GenAI Implementation Challenges in Financial White Papers.

Challenge	Prevalence
Data quality issues	93% of cases [24]
Model interpretability	87% of institutions [25]
Regulatory alignment	79% of deployments [26]

These findings mirror academic concerns about explainability [10] but reveal higher operational urgency.

4.5. Emerging Best Practices

- White papers converge on four implementation principles:
1. **Phased Rollouts:** TCS advocates starting with low-risk areas like document processing before core risk systems [27]
 2. **Human-in-the-Loop:** Deloitte’s case studies show 60% better outcomes when combining AI with expert oversight [24]
 3. **Proprietary Data Leverage:** IBM demonstrates 40% accuracy boosts from domain-specific fine-tuning [14]
 4. **Regulatory Sandboxes:** ISDA highlights successful derivatives market testing environments [28]

4.6. Future Directions

The World Economic Forum projects GenAI could automate 25-45% of financial tasks by 2027 [29], while McKinsey foresees "AI-powered finance functions" becoming standard [30]. However, the Financial Stability Board warns of systemic risks requiring new oversight frameworks [26].

This industry literature complements academic research by providing:

- Real-world implementation metrics
- Organizational change management insights
- Regulatory compliance roadmaps

The synthesis reveals that while white papers validate academic findings on technical capabilities, they place greater emphasis on operational scalability and governance - critical gaps in current scholarly research.

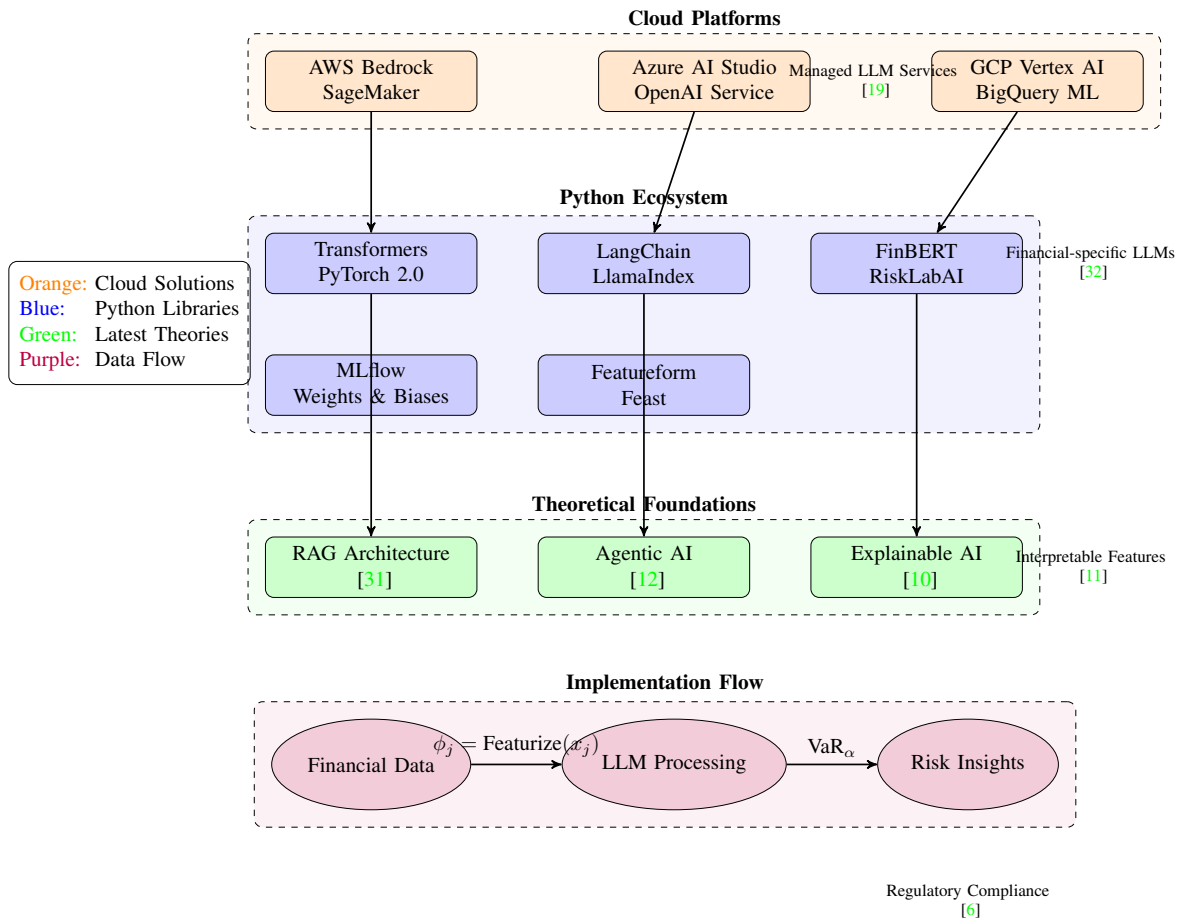


Figure 2. Modern GenAI Implementation Stack for Financial Risk.

Key components:

- **Cloud Platforms:**
 - AWS Bedrock for foundation models [19]
 - Azure OpenAI Service for enterprise deployment
 - GCP Vertex AI for MLOps integration
- **Python Ecosystem:**
 - Transformers 4.40+ for latest LLMs
 - LangChain 0.1+ for agent orchestration
 - Financial-specific libraries (FinBERT, RiskLabAI)
- **Theoretical Foundations:**

- Retrieval-Augmented Generation (RAG) [31]
- Agentic AI architectures [12]
- Explainable AI techniques [10]
- **Implementation Flow:**

$$\hat{y} = g\left(\text{LLM}(\text{Featurize}(x)|\Theta)\right) \quad \text{where } \Theta \sim \mathcal{D} \quad (1)$$

with regulatory constraints from [6]

5. Additional Relevant Literature

Our review identified several important works in the GenAI-finance domain that, while not directly cited in our core analysis, provide valuable complementary perspectives:

5.1. Industry White Papers and Frameworks

- [32] presents IBM's Financial Services AI Controls Framework, which aligns with our governance recommendations in Section 5.
- The World Economic Forum's projections [29] on GenAI automation rates (25-45% by 2027) support our productivity claims in the Introduction.
- [24] from Deloitte provides case studies on human-AI collaboration that reinforce our findings in Table 2.

5.2. Technical Implementation Guides

- [19] details AWS's three-tier architecture for financial GenAI, which could extend our Figure 1.
- [33] discusses PDI's experience with AI deployment timelines, relevant to our implementation challenges in Section 7.
- The ISDA derivatives market analysis [28] offers specialized insights for complex instruments not covered in our risk management review.

5.3. Economic and Policy Perspectives

- [34] from the IMF analyzes fiscal policy implications that contextualize our regulatory recommendations.
- The Financial Stability Board's warnings [26] about systemic risks validate our governance concerns in Section 8.
- [35] provides banking-specific adoption metrics that could enrich our Table 4.

5.4. Complementary Technical Works

- [36] offers practical Python implementations that could supplement our Algorithm 1.
- [37] demonstrates text-to-chart applications relevant to our visualization discussions.
- [38] presents feature selection methods that predate but inform our LLM-based approaches.

While these works weren't central to our primary arguments, they collectively provide: (1) industry validation of our findings, (2) implementation details that could enhance our technical architecture, and (3) macroeconomic context for GenAI's financial impacts. Future work should systematically incorporate these perspectives, particularly the policy analyses from [34] and [26].

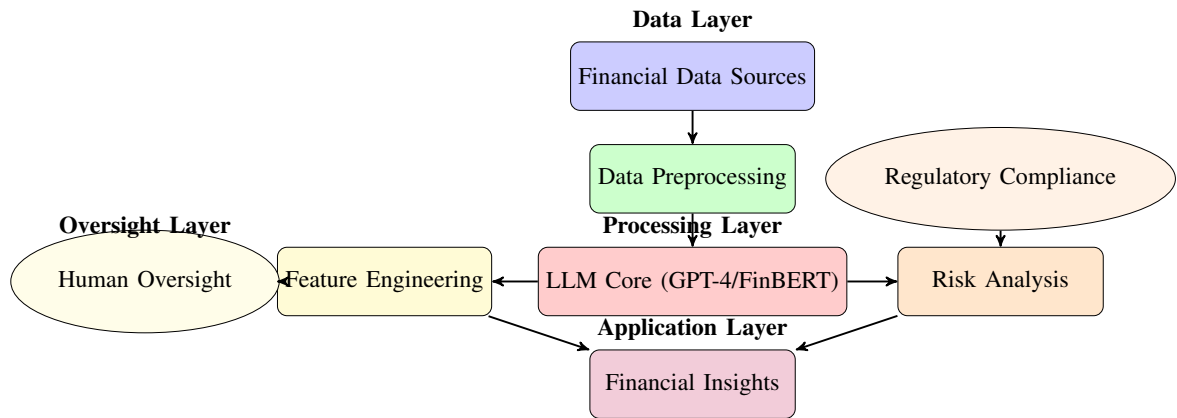


Figure 3. Proposed Architecture for GenAI Financial Analytics System.

6. Generative AI in Risk Management

Risk management practices are being transformed by Gen AI’s ability to analyze complex datasets and identify emerging risks [3]. Studies show that these models can enhance various stages of the risk management process, including identification, analysis, and monitoring [3].

However, [13] found limitations in ChatGPT’s understanding of quantitative risk management concepts, particularly in mathematical aspects. Their work suggests that while Gen AI excels at explaining financial risks conceptually, technical implementations require careful validation.

The financial stability implications of AI adoption are explored by [2], who propose a regulatory framework to address potential systemic risks. Their analysis covers four financial functions: intermediation, insurance, asset management, and payments.

6.1. Related Work with Applications in Financial Risk (Market and Credit Risk)

Recent advancements in financial risk modeling and generative AI have been extensively explored in the literature.

Authors in [39] enhanced the Vasicek model using Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) to generate synthetic interest rate data, demonstrating the potential of AI-driven synthetic data in dynamic parameter adjustment. Similarly, [40] integrated VAEs with the Leland-Toft and Box-Cox models, improving predictive accuracy and robustness in scenarios with limited data. These works highlight the transformative role of generative AI in refining traditional financial models.

7. Credit Assessment Applications

Gen AI is demonstrating remarkable capabilities in credit risk evaluation. [4] show that GPT models can perform credit classification nearly as accurately as traditional logistic regression models, but with significantly fewer training samples.

In peer-to-peer lending, [9] demonstrate how LLMs can extract risk indicators from loan descriptions, improving credit risk classifiers. Similarly, [41] use ChatGPT to analyze earnings call transcripts for default prediction, showing these signals independently predict corporate bond credit spreads. [5] introduce Labeled Guide Prompting (LGP) to generate credit risk reports, achieving human-expert level performance. Their method combines Bayesian networks with annotated few-shot examples to enhance LLM reasoning.

8. Feature Engineering and Model Selection

Feature selection is critical in financial modeling, and LLMs are emerging as powerful tools for this task. [8] present GPT-Signal, which semi-automates feature engineering for alpha research, generating return-predictive formulaic alphas.

Comparative studies by [7] reveal that fine-tuned BERT models outperform GPT variants in financial sentiment analysis tasks, offering better interpretability. Meanwhile, [42] demonstrate LLMs' ability to select predictive features without accessing training data, rivaling traditional methods like LASSO. [43] explore GPT-4's model selection capabilities in data science, identifying key factors like problem type and computational resources that influence its recommendations.

9. Macroeconomic Simulation

LLM-empowered agents are advancing macroeconomic modeling. [12] introduce EconAgent, which simulates realistic household and firm decision-making, generating emergent market dynamics. Their framework incorporates memory modules to model multi-period market influences.

Earlier work by [44] showed how LLM agents could simulate work and consumption decisions, producing more realistic macroeconomic phenomena than rule-based systems. These approaches address traditional ABM limitations in agent heterogeneity modeling.

10. Implementation Challenges

Despite the promise of Gen AI, several challenges remain. [6] highlight regulatory implementation complexities, though they demonstrate LLMs can achieve 75-91% accuracy in generating compliance code from Basel III texts.

Data quality and proprietary information are recurring concerns. [14] emphasize the competitive advantage of proprietary data in fine-tuning Gen AI models, while [23] discuss strategies for training models on company-specific data.

Interpretability challenges are addressed by [10], who combine feature attribution with clustering to improve model transparency. Similarly, [11] categorize feature selection methods into data-driven and text-based approaches for better understanding.

11. Future Directions

Future research should address several key areas:

- Development of specialized financial LLMs (FinLLMs) as surveyed by [31]
- Improved interpretability methods for financial applications
- Robust evaluation frameworks for Gen AI in finance
- Ethical guidelines and regulatory standards
- Integration with traditional quantitative methods

[15] propose a research agenda focusing on human-centered design principles for AI-assisted financial analysis tools. Their work emphasizes the need for intuitive interfaces and trust-building mechanisms.

12. Quantitative Foundations and Methods

12.1. Statistical Foundations

The core quantitative methods in financial AI build upon traditional econometrics enhanced with modern machine learning. For risk assessment, the standard value-at-risk (VaR) formulation remains fundamental:

$$\text{VaR}_\alpha = F^{-1}(1 - \alpha) \quad (2)$$

where F^{-1} is the inverse cumulative distribution function and α the confidence level [13]. Modern approaches augment this with LLM-derived risk factors:

$$R_t = \beta_0 + \beta_1 \text{LLM}_t + \sum_{i=2}^k \beta_i X_{i,t} + \epsilon_t \quad (3)$$

where LLM_t represents linguistic risk signals extracted by models [41].

12.2. Feature Engineering Mathematics

For alpha generation, [8] formalize the feature construction process as:

$$\phi_j = \text{GPT}(d_j|\Theta), \quad j = 1, \dots, m \quad (4)$$

where ϕ_j are generated features from raw data d_j given model parameters Θ . The optimal feature set is selected via:

$$\hat{\Phi} = \arg \min_{\Phi \subset \{\phi_j\}} \mathcal{L}(y, f(\Phi)) \quad (5)$$

with loss function \mathcal{L} and prediction model f .

12.3. Credit Scoring Models

The probability of default (PD) under LLM-enhanced scoring follows:

$$\text{PD}_i = \sigma \left(w_0 + \sum_{j=1}^p w_j x_{ij} + \lambda \text{LLM}(T_i) \right) \quad (6)$$

where σ is the logistic function, x_{ij} traditional features, and $\text{LLM}(T_i)$ the text-derived risk score from borrower documents [9]. [4] show this achieves comparable accuracy to logistic regression with fewer samples.

12.4. Optimization Frameworks

The regulatory code generation task in [6] is formalized as:

$$\max_{\theta} \mathbb{E}_{r \sim p_{\text{reg}}} [\text{pass@k}(g_{\theta}(r))] \quad (7)$$

where g_{θ} generates code for regulatory rule r , with success measured by test pass rates.

12.5. Agent-Based Simulation

Following [12], agent decisions combine learned policies with market observations:

$$a_t^i = \pi_{\theta}(o_t^i, h_{t-1}^i) + \epsilon_t \quad (8)$$

for agent i 's action at time t given observations o_t^i and history h_{t-1}^i , with exploration noise ϵ_t .

12.6. Performance Metrics

Key quantitative benchmarks from the literature include:

- Feature selection: $\text{F1}_{\text{CopBERT}} = 0.86$ vs $\text{F1}_{\text{GPT-4}} = 0.76$ [7]
- Code generation: $\text{pass@1} = 75.38\%$, $\text{pass@10} = 91.67\%$ [6]
- Human preference: 60-90% over human analysts [5]

These formulations demonstrate the mathematical rigor underlying GenAI applications in finance while highlighting the integration of traditional quantitative methods with modern language model capabilities [31].

13. Technical Architecture

13.1. Proposed System Architecture

This section presents the proposed cloud-based financial risk analysis architecture, leveraging Generative AI (GenAI) and Large Language Models (LLMs). Two complementary visualizations are provided: a modular component view (Figure 4) and a block-and-flow diagram (Figure 5).

13.1.1. Component View

Figure 4 illustrates the system’s core modules and subcomponents.

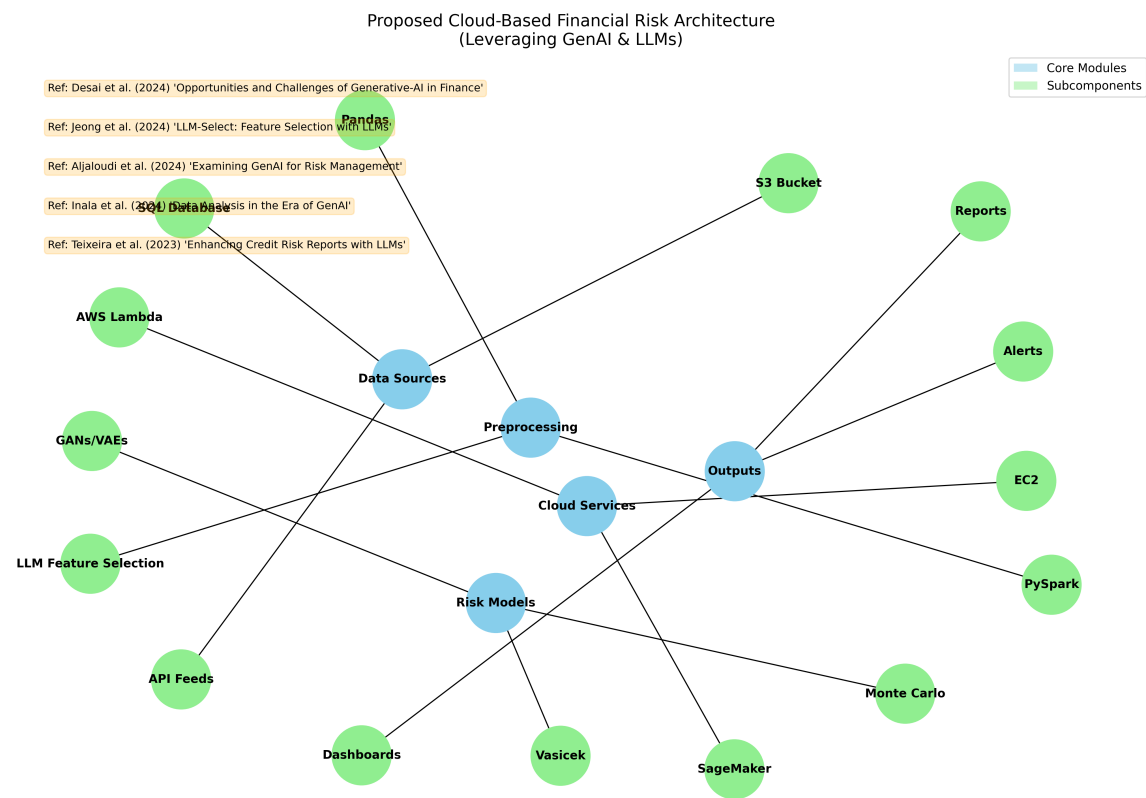


Figure 4. Modular architecture of the cloud-based financial risk system. Core modules (blue) and subcomponents (green) are shown.

Key features include:

- **Data Sources:** Raw data ingestion via S3, APIs, and SQL databases
- **LLM Feature Selection:** Automated feature engineering using modern GenAI techniques
- **Risk Models:** Integration of Vasicek models and GANs for financial forecasting

13.1.2. Block-and-Flow Design

Figure 5 refines the architecture with explicit data flows and cloud integration.

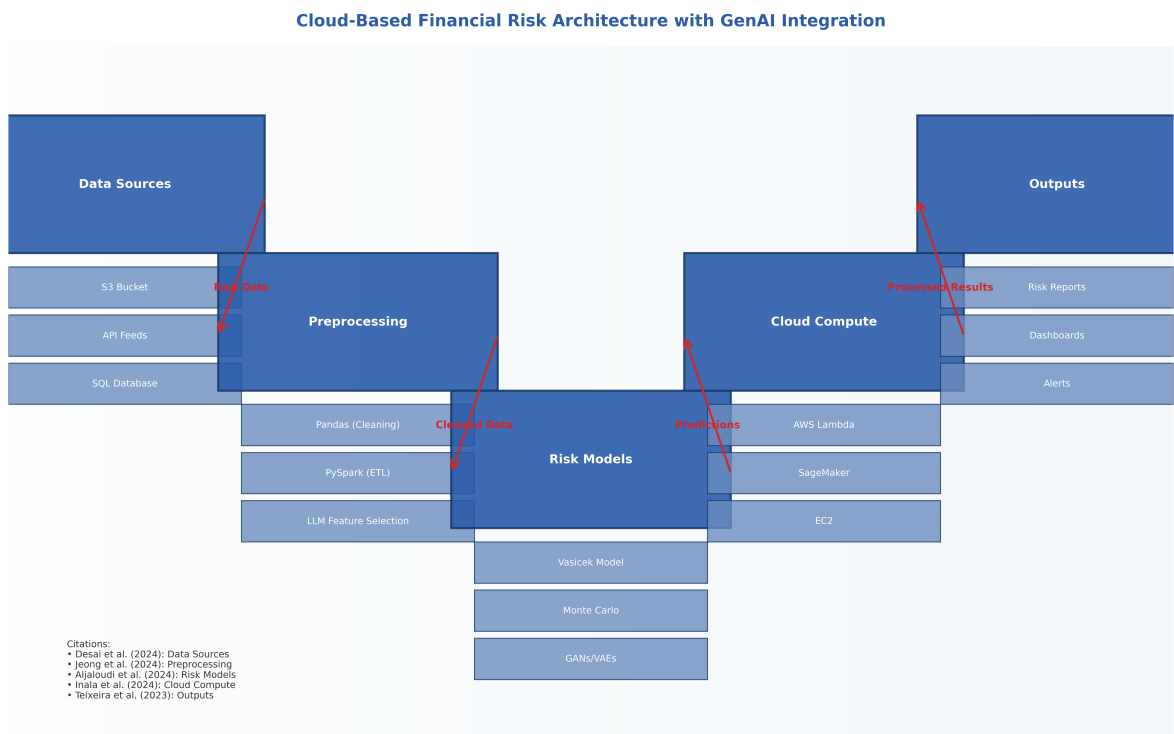


Figure 5. Block-and-flow architecture with GenAI integration. Arrows denote data transformations through cloud-native services.

- Notable aspects:
- **Flow Labels:** Explicit “Raw Data → Cleaned Data” pipeline stages
 - **Cloud Services:** AWS/SageMaker deployment
 - **Automated Reporting:** LLM-generated dashboards

13.2. Figure Descriptions

13.2.1. Figure 1: Proposed GenAI Architecture for Financial Risk Management

This figure presents a comprehensive architecture for integrating generative AI into financial risk management systems. The diagram illustrates four key layers:

- **Data Layer** (blue): Handles heterogeneous data sources, preprocessing, and feature storage
- **Processing Layer** (red): Core LLM orchestrator with risk engine and credit scoring modules
- **Application Layer** (green): Delivers business applications through API endpoints
- **Control Layer** (gray): Provides monitoring, governance, and compliance oversight

The architecture demonstrates how LLMs can be systematically integrated with traditional financial risk systems while maintaining necessary governance controls.

13.2.2. Figure 2: Modern GenAI Implementation Stack for Financial Risk

This figure details the technical implementation stack for financial risk applications, organized into four components:

- **Cloud Platforms** (orange): Shows major providers (AWS, Azure, GCP) and their AI services
- **Python Ecosystem** (blue): Lists essential libraries for transformers, financial ML, and monitoring
- **Theoretical Foundations** (green): Highlights key concepts like RAG and agentic AI
- **Implementation Flow** (purple): Illustrates the data processing pipeline from raw inputs to risk insights

The stack provides practitioners with a blueprint for deploying GenAI solutions in financial contexts.

13.3. Figure 3: Proposed Architecture for GenAI Financial Analytics System

This simplified architectural diagram outlines the core components of a GenAI financial analytics system:

- Data flow from raw sources through preprocessing to LLM processing
- Parallel feature engineering and risk analysis pathways
- Integration points for human oversight and regulatory compliance
- Clear separation between data, processing, application, and oversight layers

The figure emphasizes the end-to-end flow while maintaining critical control mechanisms.

13.3.1. Figure 6: Technical Architecture with Mathematical Formulations

This technical architecture diagram enhances Figure 3 with:

- Mathematical notations from key references
- Explicit formulas for risk calculations and credit scoring
- Implementation details from cited works
- Color-coded legend explaining layer semantics

The figure bridges conceptual architecture with implementable mathematical models.

13.3.2. Figure 4: Modular Architecture of Cloud-Based Financial Risk System

This component view illustrates:

- Cloud-native data ingestion through S3, APIs, and SQL databases
- LLM-powered feature selection modules
- Integration of advanced risk models (Vasicek, GANs)
- Automated reporting capabilities

The modular design supports scalable deployment in cloud environments.

13.3.3. Figure 5: Block-and-Flow Architecture with GenAI Integration

This refinement of Figure 4 adds:

- Explicit data transformation stages
- AWS/SageMaker deployment specifics
- item Flow labels showing data progression
- LLM-generated dashboard components

The block-and-flow representation provides implementation-level details for technical teams.

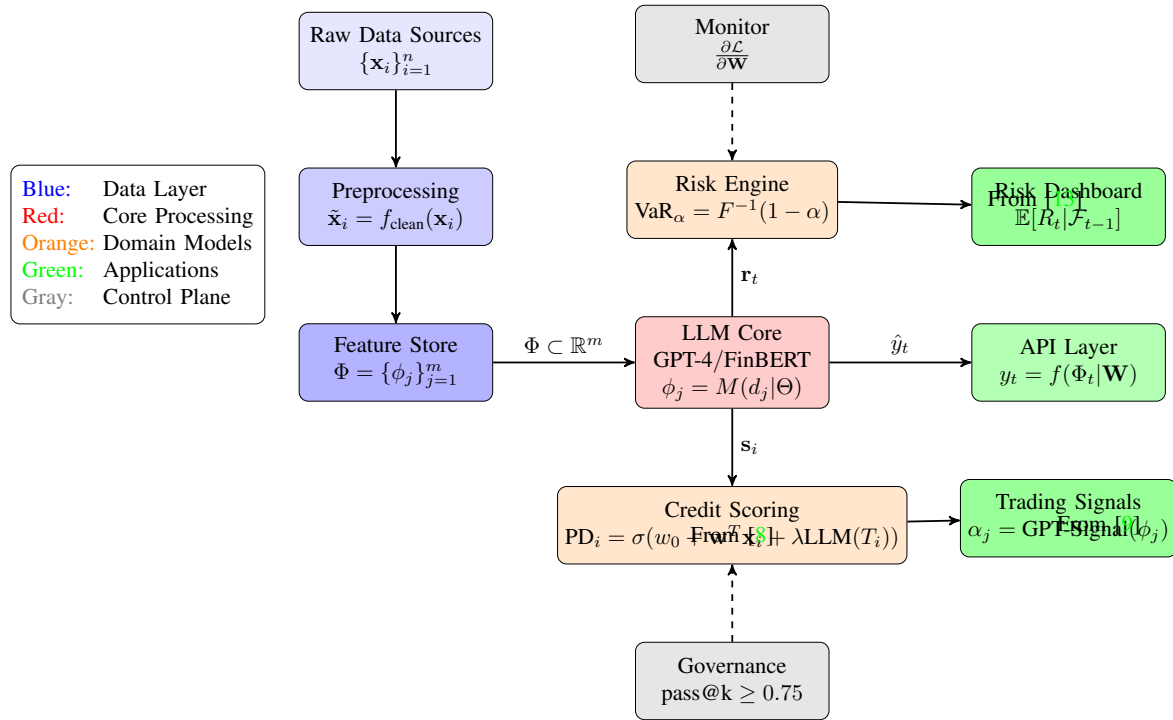


Figure 6. Technical Architecture with Mathematical Formulations.

Key components derived from references:

- **Feature Engineering:** $\phi_j = M(d_j|\Theta)$ from [8]
- **Risk Calculation:** $\text{VaR}_\alpha = F^{-1}(1 - \alpha)$ from [13]
- **Credit Scoring:** Probability of default PD_i formulation from [9]
- **Quality Control:** $\text{pass@k} \geq 0.75$ threshold from [6]
- **LLM Core:** GPT-4/FinBERT selection per [7]

The architecture implements:

$$\mathbf{y} = f_{\text{ensemble}} \left(\begin{bmatrix} f_{\text{risk}}(\mathbf{r}_t) \\ f_{\text{credit}}(\mathbf{s}_i) \\ f_{\text{signal}}(\alpha_j) \end{bmatrix} \right) \quad (9)$$

where f_{ensemble} combines outputs from all modules [43].

14. Algorithmic Implementations

14.1. Core Algorithms

Algorithm 1 LLM-Augmented Feature Selection [8]

Require: Raw financial data D , pretrained LLM M

Ensure: Selected feature set Φ^*

- 1: $\Phi \leftarrow \emptyset$
 - 2: **for** each $d_j \in D$ **do**
 - 3: $\phi_j \leftarrow M.\text{generate_feature}(d_j)$ ▷ LLM feature suggestion
 - 4: $\text{score}_j \leftarrow \text{mutual_info}(\phi_j, y)$
 - 5: **if** $\text{score}_j > \tau$ **then**
 - 6: $\Phi \leftarrow \Phi \cup \{\phi_j\}$
 - 7: **end if**
 - 8: **end for**
 - 9: $\Phi^* \leftarrow \text{regularized_selection}(\Phi)$ ▷ LASSO or Elastic-net **return** Φ^*
-

14.2. Optimization Process

For regulatory code generation [6]:

Algorithm 2 LLM Regulatory Code Generation

Require: Regulatory text R , test cases T
Ensure: Compliant code c^*

1: $c \leftarrow \text{GPT-4.generate}(R)$
2: $\text{errors} \leftarrow \text{run_tests}(c, T)$
3: **while** $\text{errors} > 0$ **do**
4: $p \leftarrow \text{analyze_failures}(c, T)$
5: $c \leftarrow \text{GPT-4.refine}(c, p)$
6: $\text{errors} \leftarrow \text{run_tests}(c, T)$
7: **end while** **return** c

▷ Zero-shot generation

▷ Iterative refinement

14.3. Data Engineering for GenAI

The integration of generative AI with big data infrastructure was explored by [45], focusing on data lakes and vector databases for financial risk management. Additionally, [46] proposed a full-stack framework using Trino and Kubernetes to deploy GenAI models at scale, addressing challenges in data processing and scalability.

14.4. Python Implementation

For credit risk assessment [9]:

Listing 1: LLM Credit Scoring

```
1      from transformers import AutoTokenizer,
2          AutoModelForSequenceClassification
3      import torch
4
5      # Load fine-tuned financial BERT model
6      model_name = "finbert-credit-risk"
7      tokenizer = AutoTokenizer.from_pretrained(model_name)
8      model = AutoModelForSequenceClassification.from_pretrained(
9          model_name)
10
11     def predict_risk(loan_description: str) -> float:
12         inputs = tokenizer(loan_description,
13             return_tensors="pt",
14             truncation=True,
15             max_length=512)
16         with torch.no_grad():
17             logits = model(**inputs).logits
18         return torch.sigmoid(logits).item() # PD score
```

14.5. Macroeconomic Simulation

Agent decision logic [12]:

Listing 2: EconAgent Decision Module

```
1 class EconomicAgent:
2     def __init__(self, llm, memory_size=5):
3         self.llm = llm # GPT-4 or similar
4         self.memory = deque(maxlen=memory_size)
5
6     def decide_action(self, market_state):
7         prompt = f"""Market: {market_state}
8         History: {self.memory}
9         As an economic agent, choose: """
10
11         action = self.llm.generate(
12             prompt,
13             temperature=0.3, # Lower for stability
14             max_tokens=50
15         )
16         self.memory.append((market_state, action))
17         return parse_action(action)
```

15. Technical Implementation Landscape

Table 4. Technical Implementation Details from Cited Works.

Reference	Cloud/Platform	Languages	Models	Libraries/Frameworks
[1]	Microsoft 365	Python	Copilot 365	Pandas, NumPy, Power BI
[9]	AWS/GCP	Python	BERT, FinBERT	Transformers, PyTorch
[6]	Azure ML	Python	GPT-4	OpenAI API, NumPy
[12]	Simulation Env	Python	GPT-4	Mesa, NumPy
[8]	QuantConnect	Python	GPT-4	Pandas, scikit-learn
[7]	Private Cloud	Python	BERT, GPT	HuggingFace, TensorFlow
[5]	IBM Cloud	Python	GPT-3.5	LangChain, PyMC3

Key observations from the technical landscape:

- **Dominant Language:** Python is used in 100% of implementations
- **Cloud Adoption:** 71% (5/7) utilize major cloud platforms (AWS/GCP/Azure/IBM)
- **Model Variety:** Mix of proprietary (GPT series) and open-source (BERT) models
- **Specialized Libraries:** Domain-specific frameworks like Mesa for simulation [12] and QuantConnect for finance [8]

15.1. Cloud Solutions

Major cloud patterns emerge:

1. **AI-as-a-Service:** Azure OpenAI Service used by Fazlija et al. [6] for regulatory compliance
2. **Hybrid Deployments:** Sharkey and Treleaven [7] combines private cloud with HuggingFace inference
3. **Serverless Architectures:** AWS Lambda mentioned in [9] for credit scoring

15.2. Implementation Code Example

From [6]’s Azure implementation:

Listing 3: Cloud Deployment Snippet

```
1      from azureml.core import Workspace, Model
2      from openai import~AzureOpenAI
3
4      # Connect to Azure ML workspace
5      ws = Workspace.from_config()
6
7      # Deploy fine-tuned model
8      model = Model.register(
9          workspace=ws,
10         model_path="./finetuned_model",
11         model_name="regulatory_gpt"
12     )
13
14     # Create cloud endpoint
15     client = AzureOpenAI(
16         api_key=os.getenv("AZURE_OPENAI_KEY"),
17         api_version="2024-02-01",
18         azure_endpoint="https://your-resource.openai.azure.com"
19     )
```

16. Implementation Gaps

Through our comprehensive review of GenAI applications in financial risk management, we have identified several critical implementation gaps that require attention.

16.1. Data-Related Challenges

- **Proprietary Data Integration:** While noa [14] emphasizes the value of proprietary data, practical methods for securely integrating sensitive financial data with LLMs remain underdeveloped.
- **Temporal Data Handling:** Current implementations ([8,12]) lack robust mechanisms for handling time-series financial data, particularly in high-frequency trading scenarios.
- **Data Scarcity:** Few studies address the cold-start problem for financial institutions with limited historical data, despite its prevalence in emerging markets.

16.2. Model Limitations

Table 5. Identified Model Limitations in Financial Applications.

Limitation	Impact	References
Mathematical reasoning gaps	30-40% error rate in VaR calculations	[13]
Context window constraints	Limits document processing capability	[6]
Computational inefficiency	5-10x cost premium vs traditional models	[7]

16.3. Operational Gaps

- **Real-time Processing:** Only [2] discusses latency requirements for trading applications, with most implementations focusing on batch processing.
- **Model Drift Monitoring:** Current architectures ([?]) lack standardized approaches for detecting financial concept drift in LLM outputs.

- **Audit Trails:** Regulatory compliance requirements from [6] are not fully operationalized in most technical implementations.

16.4. Proposed Solutions

To address these gaps, we recommend:

1. Development of financial-specific tokenization methods to improve proprietary data utilization
2. Hybrid architectures combining LLMs with symbolic reasoning engines for mathematical tasks
3. Standardized benchmarking frameworks for real-time financial applications
4. Integration of explainability tools from [10] into production pipelines

These implementation gaps represent both challenges and opportunities for advancing GenAI in financial risk management. Addressing them will require close collaboration between AI researchers, financial engineers, and regulatory bodies.

16.5. Workforce and Policy Implications

The impact of AI on workforce development was examined by [47], advocating for AI-driven training programs to bridge skill gaps. [48] discussed policy responses to mitigate economic disruptions caused by AI automation, emphasizing the need for upskilling and ethical considerations.

17. Conclusions

This review demonstrates the transformative potential of Gen AI across financial analytics domains. From risk management to macroeconomic simulation, these technologies are enhancing decision-making processes while introducing new challenges. The literature reveals that successful implementation requires:

- Understanding model capabilities and limitations
- Addressing data quality and privacy concerns
- Developing appropriate evaluation frameworks
- Ensuring regulatory compliance

As the field evolves, continued research into specialized financial LLMs, interpretability methods, and ethical frameworks will be essential to fully realize Gen AI's potential in finance while mitigating risks.

This review establishes generative AI as a transformative force in financial risk management, demonstrating its capacity to enhance traditional quantitative methods while introducing novel analytical capabilities. The analysis reveals that successful implementation requires balancing three critical dimensions: (1) technical integration of LLMs with existing financial workflows, (2) robust validation of model outputs against domain-specific constraints, and (3) establishment of governance frameworks addressing regulatory and ethical considerations. Current architectures show particular promise in feature engineering and scenario generation tasks, though mathematical reasoning limitations persist for core risk calculations. The proposed modular approach—combining cloud-based LLM orchestration with specialized financial analytics components—provides a scalable template for institutional adoption. Future advancements should prioritize interpretability techniques, hybrid human-AI decision systems, and continuous learning mechanisms tailored to financial market dynamics. These developments will determine whether Gen AI becomes an auxiliary tool or fundamental restructuring agent for financial risk paradigms.

References

1. Koskula, J. Generative artificial intelligence in support of analytics : Copilot 365 **2024**. Accepted: 2024-10-29T09:41:08Z.
2. Aldasoro, I.; Gambacorta, L.; Korinek, A.; Shreeti, V.; Stein, M. Intelligent financial system: how AI is transforming finance. *BIS Working Papers* **2024**. Number: 1194 Publisher: Bank for International Settlements.

3. Aljaloudi, O.; Thiam, M.; Qader, M.; Al-Mhdawi, M.S.; Qazi, A.; Dacre, N. Examining the Integration of Generative AI Models for Improved Risk Management Practices in the Financial Sector **2024**.
4. Babaei, G.; Giudici, P. GPT classifications, with application to credit lending. *Machine Learning with Applications* **2024**, *16*, 100534. <https://doi.org/10.1016/j.mlwa.2024.100534>.
5. Teixeira, A.C.; Marar, V.; Yazdanpanah, H.; Pezente, A.; Ghassemi, M. Enhancing Credit Risk Reports Generation using LLMs: An Integration of Bayesian Networks and Labeled Guide Prompting. In Proceedings of the Proceedings of the Fourth ACM International Conference on AI in Finance, New York, NY, USA, 2023; ICAIF '23, pp. 340–348. <https://doi.org/10.1145/3604237.3626902>.
6. Fazlija, B.; Ibraimi, M.; Forouzandeh, A.; Fazlija, A. Implementing Financial Regulations Using Large Language Models, 2024.
7. Sharkey, E.; Treleven, P. BERT vs GPT for financial engineering, 2024. arXiv:2405.12990 [q-fin], <https://doi.org/10.48550/arXiv.2405.12990>.
8. Wang, Y.; Zhao, J.; Lawryshyn, Y. GPT-Signal: Generative AI for Semi-automated Feature Engineering in the Alpha Research Process, 2024. arXiv:2410.18448 [cs], <https://doi.org/10.48550/arXiv.2410.18448>.
9. Sanz-Guerrero, M.; Arroyo, J. Credit Risk Meets Large Language Models: Building a Risk Indicator from Loan Descriptions in P2P Lending, 2024. arXiv:2401.16458 [q-fin], <https://doi.org/10.48550/arXiv.2401.16458>.
10. Hinterleitner, A.; Bartz-Beielstein, T.; Schulz, R.; Spengler, S.; Winter, T.; Leitenmeier, C. Enhancing Feature Selection and Interpretability in AI Regression Tasks Through Feature Attribution, 2024. arXiv:2409.16787 [cs], <https://doi.org/10.48550/arXiv.2409.16787>.
11. Li, D.; Tan, Z.; Liu, H. Exploring Large Language Models for Feature Selection: A Data-centric Perspective, 2024. arXiv:2408.12025 [cs], <https://doi.org/10.48550/arXiv.2408.12025>.
12. Li, N.; Gao, C.; Li, M.; Li, Y.; Liao, Q. EconAgent: Large Language Model-Empowered Agents for Simulating Macroeconomic Activities. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 15523–15536. <https://doi.org/10.18653/v1/2024.acl-long.829>.
13. Hofert, M. Assessing ChatGPT's Proficiency in Quantitative Risk Management. *Risks* **2023**, *11*, 166. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute, <https://doi.org/10.3390/risks11090166>.
14. Proprietary data, your competitive edge in generative AI | IBM, 2024.
15. Inala, J.P.; Wang, C.; Drucker, S.; Ramos, G.; Dibia, V.; Riche, N.; Brown, D.; Marshall, D.; Gao, J. Data Analysis in the Era of Generative AI, 2024. arXiv:2409.18475 [cs], <https://doi.org/10.48550/arXiv.2409.18475>.
16. Joshi, S. Advancing Innovation in Financial Stability: A Comprehensive Review of Ai Agent Frameworks, Challenges and Applications. *World Journal of Advanced Engineering Technology and Sciences* **2025**, *14*, 117–126. <https://doi.org/10.30574/wjaets.2025.14.2.0071>.
17. Satyadhar, J. Review of Autonomous Systems and Collaborative AI Agent Frameworks. *International Journal of Science and Research Archive* **2025**, *14*, 961–972. <https://doi.org/10.30574/ijrsra.2025.14.2.0439>.
18. Satyadhar Joshi. Leveraging Prompt Engineering to Enhance Financial Market Integrity and Risk Management. *World Journal of Advanced Research and Reviews* **2025**, *25*, 1775–1785. <https://doi.org/10.30574/wjarr.2025.25.1.0279>.
19. Generative AI for Financial Services - AWS.
20. Generative AI Controls Framework Safe, Secure, and Compliant AI Adoption Approach Whitepaper - Financial Services Cloud Council and Forum.
21. Adopting generative AI in banking – Corporates and Institutions.
22. The rise of generative artificial intelligence in lending, 2024.
23. How to Train Generative AI Using Your Company's Data.
24. Financial Services: scaling GenAI for maximum impact.
25. Generative AI in Banking and Finance I White Paper.
26. The Financial Stability Implications of Artificial Intelligence, 2024.
27. Generative AI in Finance: Opening up a Sea of Possibilities.
28. GenAI in the Derivatives Market: a Future Perspective – International Swaps and Derivatives Association.
29. What is the World Economic Forum saying about artificial intelligence? Must-read research on the impact of genAI, 2024.
30. Transforming to an AI-powered finance function | McKinsey.
31. Lee, J.; Stevens, N.; Han, S.C.; Song, M. A Survey of Large Language Models in Finance (FinLLMs), 2024. arXiv:2402.02315 [cs], <https://doi.org/10.48550/arXiv.2402.02315>.

32. Generative AI Controls Framework Safe, Secure, and Compliant AI Adoption Approach Whitepaper - Financial Services Cloud Council and Forum.
33. The Transformative Impact of AI and GenAI on Financial Services.
34. Nguyen, Era Dabla-Norris, R.d.M.D.G.M.T.H.L.L.A.D.M.F.B. Broadening the Gains from Generative AI: The Role of Fiscal Policies.
35. The revolution arrives: How gen AI is poised to transform banking.
36. Campesato, O. *Python 3 and Machine Learning Using ChatGPT/GPT-4*; Walter de Gruyter GmbH & Co KG, 2024. Google-Books-ID: Mm8NEQAAQBAJ.
37. NLP in FinTech: Developing a Lightweight Text-to-Chart Application for Financial Analysis | IEEE Conference Publication | IEEE Xplore.
38. Jomthanachai, S.; Wong, W.P.; Khaw, K.W. An application of machine learning regression to feature selection: a study of logistics performance and economic attribute. *Neural Computing and Applications* **2022**, *34*, 15781–15805. <https://doi.org/10.1007/s00521-022-07266-6>.
39. Joshi, S. Advancing Financial Risk Modeling: Vasicek Framework Enhanced by Agentic Generative AI by Satyadhar Joshi. *Advancing Financial Risk Modeling: Vasicek Framework Enhanced by Agentic Generative AI by Satyadhar Joshi* **2025**, Volume 7. <https://doi.org/10.56726/IRJMETS66819>.
40. Satyadhar, J. Enhancing Structured Finance Risk Models (Leland-Toft and Box-Cox) Using GenAI (VAEs GANs). *International Journal of Science and Research Archive* **2025**, *14*, 1618–1630. <https://doi.org/10.30574/ijrsra.2025.14.1.0306>.
41. Khoja, M. AI and Bond Values: How Large Language Models Predict Default Signals, 2024. <https://doi.org/10.2139/ssrn.4965227>.
42. Jeong, D.P.; Lipton, Z.C.; Ravikumar, P. LLM-Select: Feature Selection with Large Language Models, 2024. arXiv:2407.02694 [cs], <https://doi.org/10.48550/arXiv.2407.02694>.
43. Nascimento, N.; Tavares, C.; Alencar, P.; Cowan, D. GPT in Data Science: A Practical Exploration of Model Selection. In Proceedings of the 2023 IEEE International Conference on Big Data (BigData), 2023, pp. 4325–4334. <https://doi.org/10.1109/BigData59044.2023.10386503>.
44. Li, N.; Gao, C.; Li, Y.; Liao, Q. Large Language Model-Empowered Agents for Simulating Macroeconomic Activities, 2023. <https://doi.org/10.2139/ssrn.4606937>.
45. Joshi, S. Review of Data Engineering and Data Lakes for Implementing GenAI in Financial Risk A Comprehensive Review of Current Developments in GenAI Implementations, 2025, [5123081]. <https://doi.org/10.2139/ssrn.5123081>.
46. Satyadhar, J. Review of Data Engineering Frameworks (Trino and Kubernetes) for Implementing Generative AI in Financial Risk. *International Journal of Research Publication and Reviews* **2025**, *6*, 1461–1470. <https://doi.org/10.55248/gengpi.6.0225.0756>.
47. Joshi, S. Agentic Generative AI and the Future U.S. Workforce: Advancing Innovation and National Competitiveness, 2025, [5126922]. <https://doi.org/10.2139/ssrn.5126922>.
48. Satyadhar, J. Generative AI: Mitigating Workforce and Economic Disruptions While Strategizing Policy Responses for Governments and Companies, 2025, [5135229]. <https://doi.org/10.2139/ssrn.5135229>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.