European Journal of AI, Computing & Informatics

Vol. 2 No. 1 2026

Review Open Access

Study of Generative AI in Automated Financial Report Generation

Jialong Li [1,*]

[1] Ethic Inc., Jersey City, New Jersey, USA
* Correspondence: Jialong Li, Ethic Inc., Jersey City, New Jersey, USA

Received: 13December 2025
Revised: 01 March 2026
Accepted: 16 March 2026
Published: 19 March 2026

Abstract: This paper provides a comprehensive review of the rapidly evolving field of Generative AI in automated financial report generation. Traditionally, financial report generation has been a labor-intensive process relying on manual data aggregation, analysis, and narrative composition. However, recent advancements in Generative AI, particularly Large Language Models (LLMs), have demonstrated the capacity to automate and significantly enhance this process. This paper describes the current state-of-the-art, tracing the historical development of AI techniques applied to financial reporting. The paper examines core themes such as the application of LLMs for narrative generation from structured financial data, and the use of Generative Adversarial Networks (GANs) for synthetic data generation and fraud detection. A critical comparison of different Generative AI models is presented, highlighting their strengths and weaknesses in the context of financial reporting, alongside a discussion of the inherent challenges, including data bias, regulatory compliance, and the need for explainable AI. Finally, the paper explores future research directions, such as the integration of multi-modal data, the development of more robust and transparent AI models, and the ethical considerations surrounding the widespread adoption of Generative AI in finance. This review aims to provide researchers, practitioners, and regulators with a thorough understanding of the opportunities and challenges presented by Generative AI in transforming the landscape of financial report generation.

Keywords: generative AI; financial report generation; large language models; automated reporting; financial technology; natural language processing; financial data analysis

1. Introduction

1.1. Motivation and Background

Financial reports are crucial for stakeholders to understand a company's performance and financial health. These reports, including balance sheets, income statements, and cash flow statements, inform investment decisions and regulatory compliance. Traditionally, their creation involves manual data collection, consolidation, and analysis, a process that is often time-consuming, error-prone, and resource-intensive. Automating this process using Generative AI offers the potential to significantly improve efficiency, reduce errors, and provide more timely and insightful financial information. The use of *AI* can also allow for more dynamic and personalized reports [1].

1.2. Scope and Objectives

This review focuses on the application of generative AI in automating financial report generation. The scope encompasses Large Language Models (LLMs) and other generative techniques used for processing financial data, extracting key insights, and producing narrative reports. Our objectives are threefold: first, to provide a comprehensive overview of current generative AI methodologies employed in this domain; second, to identify the key challenges, such as data privacy (x), accuracy (y), and regulatory compliance (z); and third, to explore potential future research directions and opportunities for improvement in automated financial reporting [2].

2. Historical Overview of AI in Financial Reporting

2.1. Early Rule-Based Systems

Early attempts at automating financial reporting leveraged rule-based systems and expert systems. These systems encoded accounting principles and reporting standards as a series of "if-then" rules. For example, "IF $revenue > \$1{,}000{,}000$ THEN classify as large enterprise." Expert systems aimed to capture the knowledge of experienced accountants to guide report generation. However, these early systems suffered from limitations. They were brittle, requiring extensive manual updates to reflect evolving regulations. Furthermore, they struggled with unstructured data and lacked the ability to handle nuanced judgments inherent in financial reporting, ultimately proving inflexible and difficult to scale. The characteristics and limitations of these early systems, compared to more modern approaches, are summarized in Table 1.

Table 1. Evolution of AI Techniques in Financial Reporting.

Stage	Technique	Description	Limitations
Early Automation	Rule-Based Systems	Encoded accounting principles as "if-then" rules (e.g., "IF $revenue > \$1{,}000{,}000$ THEN classify as large enterprise").	Brittle, required manual updates for regulatory changes, struggled with unstructured data, and lacked nuanced judgment.
Early Automation	Expert Systems	Captured the knowledge of experienced accountants to guide report generation.	Inflexible, difficult to scale, and unable to handle complex or changing financial landscapes.

2.2. Statistical and Machine Learning Approaches

Statistical models have long been employed in financial reporting for tasks like fraud detection and risk assessment. Regression analysis, for instance, can identify anomalies by modeling expected values and flagging deviations exceeding a threshold, such as 3σ. Traditional machine learning techniques, including Support Vector Machines (SVMs) and decision trees, offer more sophisticated pattern recognition. SVMs excel at classifying transactions based on multiple features, while decision trees provide interpretable rules for identifying high-risk activities. These methods often rely on labeled datasets to train models that predict the likelihood of fraudulent behavior or assess credit risk based on factors like debt-to-income ratio (d/i).

2.3. The Rise of Deep Learning

Deep learning's emergence significantly impacted financial reporting automation. Recurrent Neural Networks (RNNs), adept at handling sequential data, found application in processing financial time series like stock prices and trading volumes, enabling forecasting and anomaly detection. Convolutional Neural Networks (CNNs), initially designed for image recognition, were adapted to analyze textual information within financial reports, such as extracting key performance indicators (KPIs) and sentiment analysis from news articles and company filings. This marked a shift towards more sophisticated feature extraction and pattern recognition compared to traditional machine learning techniques. The ability of deep learning models to learn complex relationships from high-dimensional data x_i offered a new avenue for automating report generation [3].

3. Core Theme A: LLMs for Narrative Generation

3.1. LLMs for Text Summarization and Simplification

Large Language Models (LLMs) offer significant capabilities for enhancing financial report generation through text summarization and simplification. The inherent complexity of financial data often necessitates concise summaries to highlight key performance indicators and trends. LLMs can be leveraged to distill lengthy reports into shorter, more digestible formats, improving information accessibility for stakeholders. Furthermore, LLMs can simplify the technical jargon prevalent in financial documents. By rephrasing complex sentences and replacing specialized terminology with more common language, LLMs can broaden the audience capable of understanding the report's content. This simplification process can be guided by parameters such as target reading level, measured by metrics like the Flesch-Kincaid grade level, denoted as F. The goal is to minimize F while preserving the report's accuracy and integrity [4].

3.2. Generating Financial Narratives from Structured Data

LLMs offer a robust framework for automating the generation of financial narratives from structured data. This process involves transforming tables of figures, such as balance sheets and income statements, into human-readable text. The LLM acts as a translator, identifying key trends and relationships within the data. For example, a decrease in accounts receivable, represented as A_r, coupled with an increase in sales, S, could be narrated as improved efficiency in collecting payments. Furthermore, LLMs can contextualize these trends by comparing them to previous periods or industry benchmarks. A critical challenge lies in ensuring accuracy and avoiding misinterpretations of the underlying financial data. Careful prompt engineering and validation are crucial for reliable narrative generation [5]. Table 2 provides a practical example of a narrative generated by an LLM based on specific financial input data.

Table 2. Example of LLM-Generated Narrative from Financial Data.

Financial Data	LLM-Generated Narrative
Decrease in Accounts Receivable (A_r) + Increase in Sales (S)	"The company demonstrates improved efficiency in collecting payments from its customers, as evidenced by a decrease in accounts receivable coupled with an increase in sales."
Increase in Net Income (NI) compared to previous year	"Net income significantly increased compared to the previous year, indicating improved profitability and operational performance."
Increase in Debt-to-Equity Ratio	"The company's debt-to-equity ratio has increased, suggesting greater reliance on debt financing compared to equity."
Gross Profit Margin declined by 5%	"The Gross Profit Margin declined by 5%, which may indicate rising costs of goods sold or pricing pressures."
Return on Equity (ROE) above industry average	"The company's Return on Equity (ROE) is performing above the industry average, signaling superior profitability compared to its peers."

3.3. Prompt Engineering for Financial Reporting

Prompt engineering significantly impacts the quality of LLM-generated financial narratives. Detailed instructions specifying the target audience, report type (e.g., 10-K, earnings call transcript), and key performance indicators ($KPIs$) generally yield more accurate and relevant results. Few-shot learning, providing examples of well-written financial reports, further refines the LLM's output style and content. Strategies include specifying desired tone (e.g., conservative, optimistic), incorporating specific financial ratios (ROA, ROE), and explicitly requesting explanations for significant variances. Iterative prompt refinement, evaluating the LLM's output and adjusting the prompt accordingly, is crucial for optimizing narrative quality [6].

4. Core Theme B: GANs & Synthetic Financial Data

4.1. GANs for Synthetic Financial Data Generation

Generative Adversarial Networks (GANs) present an effective approach for generating synthetic financial datasets. These networks, composed of a generator and a discriminator, learn to mimic the statistical properties of real financial data. The generator creates synthetic data instances, while the discriminator attempts to distinguish between real and generated samples. Through an adversarial training process, the generator improves its ability to produce realistic synthetic data that reflects the complexities of financial markets. This approach addresses the challenge of data scarcity, particularly for specialized financial instruments or events. Furthermore, GANs can be used to create privacy-preserving datasets. Since the synthetic data is not directly derived from real individuals or institutions, it can be shared and analyzed without revealing sensitive information, thus protecting privacy while enabling research and development. The quality of synthetic data is crucial; metrics like statistical similarity and utility for downstream tasks are used to evaluate GAN performance [7].

4.2. GANs for Fraud Detection

GANs provide a powerful framework for fraud detection by learning the underlying distribution of normal financial transactions. The architecture of this system, comprising a generator and a discriminator, is detailed in Table 3. The generator network, G, is trained to produce synthetic financial data that mimics legitimate transactions, while the discriminator network, D, learns to distinguish between real and generated data. Once trained, the GAN captures the complex patterns of normal financial activity. Fraudulent transactions, being anomalous, deviate significantly from this learned distribution. These anomalies are then flagged by the discriminator, which assigns a low probability to their authenticity. The anomaly score, S, can be defined as $S = 1 - D(x)$, where x represents a financial transaction. Higher S values indicate a higher likelihood of fraud, enabling proactive detection and mitigation of fraudulent activities [8].

Table 3. GAN-based Fraud Detection System Architecture.

Component	Description
Generator (G)	A neural network that learns to generate synthetic financial transactions resembling legitimate ones. Trained to fool the discriminator.
Discriminator (D)	A neural network that learns to distinguish between real and generated financial transactions. It attempts to identify fraudulent transactions as anomalies.
Input Data (x)	Represents a financial transaction represented as a vector of features, serving as input to both the generator and discriminator.
Anomaly Score (S)	Calculated as $S = 1 - D(x)$, where x is a financial transaction and $D(x)$ is the discriminator's output (probability of x being real). A higher S indicates a higher likelihood of fraud.
Training Process	The generator G and discriminator D are trained adversarially.G tries to generate realistic transactions to fool D, while D tries to correctly classify real and generated transactions.
Output	A fraud detection score (anomaly score S) for each transaction, allowing for the ranking and flagging of potentially fraudulent activities.

4.3. Addressing Data Imbalance with GANs

Data imbalance poses a significant challenge in financial applications like fraud detection, where fraudulent transactions ($y = 1$) are far less frequent than legitimate ones ($y = 0$). Generative Adversarial Networks (GANs) offer a promising solution by synthesizing examples of the minority class. A GAN, comprising a generator and a discriminator, learns to create realistic synthetic data points. The generator attempts to produce data that mimics the characteristics of the minority class, while the discriminator tries to distinguish between real and generated samples. Through adversarial training, the generator improves its ability to create convincing synthetic data, effectively augmenting the minority class and mitigating the impact of data imbalance on model performance [9].

5. Comparison of Methods and Challenges

5.1. Comparative Analysis of Generative AI Models

Different generative AI models offer varying strengths for automated financial report generation. Large Language Models (LLMs) excel in coherence and narrative generation, producing human-readable reports. However, ensuring numerical accuracy remains a challenge, often requiring post-generation verification. Generative Adversarial Networks (GANs) can be trained to mimic the statistical properties of financial data, potentially improving accuracy in specific areas like fraud detection reports. Yet, GANs often struggle with generating long, coherent documents. Variational Autoencoders (VAEs) offer a balance, enabling controlled generation through latent space manipulation. The efficiency of each model also differs; LLMs generally require more computational resources for training and inference compared to GANs or VAEs. The choice of model depends on the specific requirements of the financial report, balancing accuracy, coherence, and computational cost (C) [10].

5.2. Challenges and Limitations

Generative AI applications in financial reporting face several critical challenges, as summarized in Table 4. Data bias present in training datasets can lead to skewed or discriminatory outputs, impacting the fairness and accuracy of reports. The “black box” nature of many AI models necessitates explainable AI (XAI) to ensure transparency and auditability, crucial for stakeholder trust. Regulatory compliance, particularly with standards like Sarbanes-Oxley (SOX) and IFRS, demands rigorous validation and control mechanisms. Furthermore, the potential for generating misleading or inaccurate financial information poses a significant risk, requiring robust validation and human oversight to mitigate potential errors and ensure the reliability of generated reports. The cost of implementing and maintaining these systems can also be a barrier [11].

Table 4. Challenges in Generative AI for Financial Reporting.

Challenge	Description
Data Bias	Training datasets may contain biases, leading to skewed or discriminatory outputs in financial reports.
Lack of Explainability (Black Box)	The “black box” nature of many AI models hinders transparency and auditability, creating trust and validation problems. Explainable AI (XAI) is needed.
Regulatory Compliance	Meeting regulatory requirements like Sarbanes-Oxley (SOX) and IFRS necessitates stringent validation and control mechanisms.
Potential for Misleading Information	Generative AI may create inaccurate or misleading financial data, demanding robust validation and human oversight.
Implementation and Maintenance Costs	The cost of deploying and maintaining generative AI systems can be a significant obstacle.

6. Future Perspectives

6.1. Multi-Modal Data Integration

The future of automated financial report generation lies in the integration of multi-modal data. Current systems primarily rely on structured numerical data and textual narratives. Expanding the input to include images, such as graphs embedded in presentations or photographs of physical assets, and videos, like CEO interviews or factory tours, can significantly enrich the reports. Generative AI models, particularly those leveraging transformers, can be trained to process and synthesize these diverse data types. For example, an image recognition model could identify key assets from a photograph, while a natural language processing model summarizes the sentiment expressed in a video interview. This information can then be incorporated into the report narrative, providing a more comprehensive and engaging overview of the company's financial performance and strategic direction. The challenge lies in developing robust methods for aligning and interpreting these disparate data streams, ensuring accuracy and avoiding spurious correlations. The potential benefits, however, in terms of report diversity and richness, are substantial [12].

6.2. Towards Explainable and Transparent AI

The adoption of Generative AI in financial reporting hinges on building trust and ensuring accountability. Current models often operate as black boxes, hindering understanding of how specific outputs are derived. Future research must prioritize developing more explainable and transparent AI. Techniques like attention mechanisms can highlight the data points most influential in generating a particular report section, providing insights into the model's reasoning. Furthermore, exploring interpretability methods, such as SHAP values or LIME, can quantify the contribution of different input features (x_i) to the final output (y), thereby increasing transparency. This shift towards explainability is crucial for auditors, regulators, and stakeholders to confidently rely on AI-generated financial reports [3].

7. Conclusion

Generative AI demonstrates significant potential in automating financial report generation. Our review highlights its ability to streamline data aggregation, narrative generation, and report formatting, potentially reducing report creation time and costs. Key findings indicate that models like Large Language Models (LLMs) can effectively translate numerical data into coherent textual summaries. However, challenges remain. Ensuring accuracy, mitigating bias, and maintaining compliance with regulatory standards are critical concerns. Furthermore, the "black box" nature of some models raises questions about transparency and auditability, requiring further research into explainable AI (XAI) techniques within this domain.

Generative AI holds immense promise for automating financial report generation, potentially increasing efficiency and reducing costs. However, the current state necessitates careful consideration of accuracy and reliability. Future research should focus on enhancing model transparency and mitigating biases to ensure fairness and prevent the propagation of errors. Furthermore, as Generative AI becomes more integrated into financial reporting, addressing ethical concerns and establishing clear regulatory frameworks are paramount to maintain trust and accountability in the financial ecosystem. The responsible development and deployment of these technologies are crucial for realizing their full potential while safeguarding against potential risks.

References

1. E. Blankespoor, E. deHaan, and Q. Li, "Generative AI in financial reporting," Financial Reporting, 2025.
2. R. Modak, "Generative AI for automated business report generation and analysis," World Journal of Advanced Engineering Technology and Sciences, vol. 15, no. 2, pp. 10-30574, 2025.
3. D. K. C. Lee, C. Guan, Y. Yu, and Q. Ding, "A comprehensive review of generative AI in finance," FinTech, vol. 3, no. 3, pp. 460-478, 2024.
4. N. K. Shukla, R. Katikeri, M. Raja, G. Sivam, S. Yadav, A. Vaid, and S. Prabhakararao, "Generative AI approach to distributed summarization of financial narratives," in 2023 IEEE International Conference on Big Data (BigData), 2023, pp. 2872-2876.
5. Y. Ren, W. Hu, Z. Wang, X. Zhang, Y. Wang, and X. Wang, "A hybrid deep generative neural model for financial report generation," Knowledge-Based Systems, vol. 227, 107093, 2021.
6. A. Francis and C. X. Zhang, "The Rise of Generative AI in Financial Analysis: Evidence from Equity Research," Working Paper, 2024.
7. S. Amri, R. Bani, and S. Bani, "An approach to the analysis of financial documents using generative AI," in Proc. 7th Int. Conf. Networking, Intelligent Systems and Security, 2024, pp. 1-5.

8. H. R. Paknoodezh and S. Bagherpanah, "Designing an Intelligent Financial Reporting Framework Based on Generative Artificial Intelligence (AI-FRS): A Step Toward Self-Adaptive Accounting Standards," Business, Marketing, and Finance Open, pp. 1-17, 2025.
9. S. Agarwal, D. M. Dash, A. F. H. M. Dass, S. B. R. Pagilla, C. V. Sanaboina, M. A. Lanham, and A. Mishra, "Enhancing Financial Analysis with Generative AI: Utilizing Large Language Models for Efficient Data Extraction," in World Congress in Computer Science, Computer Engineering & Applied Computing, Cham, Springer Nature Switzerland, 2024, pp. 244-259.
10. Y. Chen, S. Yan, Y. Jin, J. Li, and X. Jin, "Generative Artificial Intelligence in Finance: A Systematic Literature Review and a Research Agenda," Digital Technologies Research and Applications, vol. 4, no. 2, pp. 14-32, 2025.
11. W. Mbanyele, "Generative AI and ChatGPT in Financial Markets and Corporate Policy: A Comprehensive Review," Working Paper, 2024.
12. M. T. Afzal, "Beyond The Bottom Line: Using Generative AI to Extract Narrative Financial Disclosures," Working Paper, 2024.