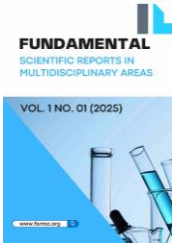
Fundamental Scientific Reports in Multidisciplinary Areas

Vol: 02 Issue:03 2026

E-ISSN: 3134-7517 P-ISSN: 3134-7509



ARTICLE HISTORY

Received: [10 October 2026]
Revised: [17 October 2026]
Accepted: [24 October 2026]

Published online: [31 October 2026]

KEYWORDS

Dynamic Risk Segmentation; Graph Neural Networks; Retrieval-Augmented Generation; Multi-Agent Large Language Models; Multimodal Financial Analysis; Financial Risk Management; Automated Investment Decision-Making

LICENSE

This article is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). You may share and adapt the material in any medium or format, with appropriate credit to the original author(s) and a link to the licence.
https://creativecommons.org/licenses/by/4.0/

COPYRIGHT

© 2026 The Author(s). Authors retain copyright; FSRMA holds a non-exclusive licence to publish.

ORIGINAL RESEARCH ARTICLE

DynFinRAG-Agent: A Dynamic Risk Segmentation and Graph-Enhanced Retrieval-Augmented Multi-Agent Framework for Multimodal Financial Risk Analysis and Automated Investment Decision-Making

Yang Shi[1], Yujie Lu[2], Shaoli Qian[3]
[1] Financial Mathematics, Case Western Reserve University, Cleveland OH, USA
2 Columbia University in the City of New York, New York, NY, USA
[3] Amazon.com Services LLC, Seattle, WA, USA
[1] yxs879@case.edu
[2] yl5266@columbia.edu
[3] qsl@terpmail.umd.edu

ABSTRACT

The increasing heterogeneity of financial data and rapidly evolving user behaviors pose significant challenges to financial risk monitoring and automated investment decision-making, particularly when static risk Aprofiles and isolated models fail to capture temporal risk transitions, relational anomalies, and knowledge-intensive financial contexts. This study proposes DynFinRAG-Agent, a Dynamic Risk Segmentation and Graph-Enhanced Retrieval-Augmented Multi-Agent framework for multimodal financial risk analysis and automated investment decision-making. The framework integrates a Multimodal Financial Encoder (MFE) to jointly represent market time-series, corporate financial indicators, user behavioral records, and financial texts, while Dynamic Risk Segmentation (DRS) continuously identifies evolving user risk states. A Dynamic Risk Graph Network (DRG-Net) further models interactions among users, accounts, firms, assets, and financial events, and its risk representations are incorporated into Graph-Enhanced Retrieval-Augmented Generation (GF-RAG) to retrieve relevant regulatory rules, historical cases, and market knowledge. Specialized Multi-Agent Large Language Models (LLMs) then perform risk assessment, compliance analysis, market intelligence extraction, portfolio evaluation, and investment decision generation. Experiments on a temporally and entity-aligned benchmark constructed from multiple public financial datasets demonstrate that DynFinRAG-Agent achieves 95.1% accuracy and an Area Under the Curve (AUC) of 0.978 for financial risk prediction, while reaching 93.5% decision consistency in automated investment decision-making. These results indicate that dynamic risk segmentation, relational graph reasoning, and evidence-grounded multi-agent collaboration jointly improve the reliability and effectiveness of multimodal financial risk analysis and investment decisions.

1. Introduction

The rapid expansion of financial markets and digital financial services has generated increasingly heterogeneous data, including market time-series, corporate financial statements, transaction records, user behavioral information, financial news, regulatory documents, and investment reports. Such data provide valuable signals for financial risk monitoring and investment decision-making, but their heterogeneous structures, temporal dynamics, and complex interdependencies make comprehensive analysis increasingly difficult. Conventional financial risk models typically rely on static customer profiles, structured financial indicators, or isolated predictive models, which limits their ability to identify continuously evolving risk states, capture relational risk propagation, and integrate external financial knowledge. These limitations are particularly significant in automated investment environments, where reliable decisions require not only accurate risk prediction but also contextual reasoning, evidence retrieval, compliance assessment, and adaptive portfolio management. Dynamic risk segmentation provides an effective mechanism for addressing the temporal nature of financial risk. Unlike static segmentation approaches that assign users or financial entities to fixed categories, dynamic risk segmentation continuously updates risk states according to changes in transaction behavior, financial conditions, market environments, and historical risk patterns. A financial entity may therefore transition from normal to potential-risk, abnormal, or high-risk states as new evidence emerges. However, conventional clustering and classification methods generally treat entities independently and cannot adequately represent the relationships among customers, accounts, firms, securities, funds, products, and financial events. Graph Neural Networks (GNNs) provide a natural solution by representing these entities and their interactions as dynamic financial graphs, enabling relational dependencies, community-level anomalies, and risk propagation patterns to be jointly modeled. Meanwhile, Retrieval-Augmented Generation (RAG) can enhance financial intelligence by grounding model reasoning in regulatory policies, historical risk cases, corporate disclosures, market information, and expert knowledge. Nevertheless, conventional RAG mainly relies on semantic similarity and may overlook the structural relationships underlying financial risk. Combining graph representations with retrieval can therefore improve the relevance and contextuality of retrieved evidence. Furthermore, Large Language Model (LLM)-based multi-agent systems can distribute complex financial tasks among specialized agents, enabling coordinated risk assessment, regulatory compliance analysis, market intelligence extraction, portfolio evaluation, and investment strategy generation. Integrating these capabilities with multimodal financial representations creates an opportunity to develop a more adaptive and evidence-grounded decision framework. To address these challenges, this study proposes DynFinRAG-Agent, a Dynamic Risk Segmentation and Graph-Enhanced Retrieval-Augmented Multi-Agent framework for Multimodal Financial Risk Analysis and Automated Investment Decision-Making. The framework first employs a Multimodal Financial Encoder to integrate market time-series, financial indicators, behavioral information, and financial text. Dynamic Risk Segmentation then identifies evolving risk states, while a Dynamic Risk Graph Network models relational dependencies and risk propagation among financial entities. The resulting risk-aware graph representations are incorporated into Graph-Enhanced Retrieval-Augmented Generation to retrieve structurally and semantically relevant financial knowledge. Finally, specialized LLM agents collaboratively perform risk assessment, compliance analysis, market interpretation, portfolio evaluation, and risk-aware investment decision generation.

The main contributions of this study are summarized as follows:

A dynamic multimodal financial risk framework is developed to jointly model heterogeneous financial information and continuously evolving risk states rather than relying on static risk profiles.

A dynamic graph-enhanced risk modeling mechanism is introduced to capture individual anomalies, relational dependencies, community-level risks, and cross-entity risk propagation.

© 2026 The Author(s). Licensed under CC BY 4.0

A graph-enhanced RAG mechanism is designed to integrate risk states and financial graph structures into knowledge retrieval, improving evidence grounding for financial reasoning and decision generation.
A collaborative multi-agent decision architecture is established in which specialized LLM agents jointly conduct risk analysis, compliance verification, market intelligence extraction, portfolio assessment, and automated investment decision-making. An integrated risk-aware investment decision framework is constructed to connect dynamic risk identification with evidence-grounded reasoning and adaptive investment strategies, providing a unified approach to reliable and explainable financial intelligence.

2.Literature Review

Recent advances in artificial intelligence have substantially transformed financial risk management, financial information processing, and automated investment analysis. Existing studies have progressively moved from conventional statistical and machine learning models toward graph-based risk modeling, domain-specific Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), and multi-agent decision systems. However, these research directions have largely been investigated independently. This section reviews the major research lines related to DynFinRAG-Agent, including financial risk and fraud detection, graph neural networks for relational financial analysis, multimodal and domain-specific financial LLMs, retrieval-augmented financial reasoning, and LLM-based multi-agent investment decision-making.

2.1 Financial Risk Analysis and Intelligent Fraud Detection

Traditional financial risk analysis has mainly relied on statistical models, rule-based systems, and conventional machine learning techniques. Although these approaches are relatively interpretable and computationally efficient, they often depend on manually engineered features and static assumptions, making them less effective when financial behaviors evolve rapidly. Hilal et al. [1] systematically reviewed financial fraud detection methods and highlighted the increasing importance of semi-supervised and unsupervised learning for identifying evolving and previously unseen fraudulent patterns. Their study also emphasized that modern financial fraud increasingly involves complex behavioral and structural relationships that cannot be adequately represented by isolated transaction-level features.
Zhu et al. [2] further reviewed intelligent financial fraud detection in the post-pandemic era and emphasized the transition from conventional quantitative transaction data toward heterogeneous data sources. Their analysis identified Graph Neural Networks (GNNs) as an important direction because financial entities and transactions naturally form interconnected structures. These studies demonstrate the limitations of treating financial entities independently and motivate the development of dynamic risk representations capable of integrating behavioral, temporal, and relational information. However, most existing financial risk studies primarily focus on fraud or anomaly detection rather than continuously updating risk segments and connecting risk identification with downstream knowledge-grounded investment decisions.

2.2 Graph Neural Networks for Relational Financial Risk Modeling

GNNs have become an important approach for financial risk analysis because financial activities can naturally be represented as graphs containing customers, accounts, firms, assets, transactions, and financial events. GraphSAGE proposed by Hamilton et al. [3] introduced an inductive graph representation framework that learns node representations through neighborhood aggregation, enabling the model to generalize to previously unseen nodes. Veličković et al. [4] proposed Graph Attention Networks (GATs), which use attention mechanisms to assign different importance to neighboring nodes and therefore provide a flexible mechanism for modeling heterogeneous relational dependencies. For dynamic financial environments, Rossi et al. [5] introduced Temporal Graph Networks (TGNs), which incorporate temporal events and memory mechanisms into graph representation learning. Such temporal graph modeling is particularly relevant to financial systems in which transaction relationships and risk states continuously evolve. Cheng et al. [6] developed a spatial-temporal attention-based GNN for credit card fraud detection

© 2026 The Author(s). Licensed under CC BY 4.0

and demonstrated that jointly modeling transaction relationships and temporal information can improve fraud detection performance. More recently, Motie and Raahemi [7] systematically reviewed GNN-based financial fraud detection and identified substantial research activity involving static, temporal, and heterogeneous financial graphs. Nevertheless, existing studies generally concentrate on fraud classification or anomaly identification. They rarely integrate dynamic risk segmentation, graph-based risk propagation, external knowledge retrieval, and downstream investment decision-making within a unified framework.

2.3 Multimodal Financial Representation and Financial Large Language Models

The increasing availability of financial reports, market data, news, charts, and regulatory documents has stimulated research on domain-specific and multimodal LLMs. FinQA, introduced by Chen et al. [8], established a benchmark for numerical reasoning over financial reports and demonstrated that financial reasoning requires the joint understanding of textual and tabular information. This finding indicates that financial intelligence systems need to move beyond purely textual modeling and incorporate heterogeneous financial evidence.

Huang et al. [9] developed FinBERT, a finance-adapted language model that incorporates domain-specific financial knowledge and contextual information for financial text analysis. BloombergGPT, proposed by Wu et al. [10], further demonstrated the effectiveness of large-scale domain-specific pretraining for financial natural language processing, covering applications such as sentiment analysis, information extraction, and question answering. Yang et al. [11] introduced FinGPT as an open-source financial LLM framework and emphasized data-centric development and lightweight adaptation for applications including robo-advising and algorithmic trading. In addition, Wang et al. [12] proposed FinVis-GPT, a multimodal LLM designed for financial chart analysis, demonstrating the potential of combining visual financial information with language-based reasoning. Although these models substantially improve financial information understanding, most existing approaches focus on individual modalities or specific financial NLP tasks. They generally do not explicitly model dynamic risk transitions or the relational structure among financial entities. DynFinRAG-Agent therefore extends multimodal financial representation toward a risk-aware setting by jointly considering market time-series, structured financial indicators, behavioral information, and financial text.

2.4 Retrieval-Augmented Generation and Graph-Enhanced Financial Retrieval

RAG has emerged as an effective approach for grounding LLM responses in external and updatable knowledge. Lewis et al. [13] introduced the original Retrieval-Augmented Generation framework, demonstrating that external document retrieval can complement the knowledge stored within pretrained language models. In financial applications, retrieval is particularly important because regulatory requirements, corporate disclosures, market events, and financial policies change continuously and cannot be reliably encoded only through model parameters.

Recent research has further explored the integration of graph structures into RAG. Peng et al. [14] provided a comprehensive survey of Graph Retrieval-Augmented Generation (GraphRAG), categorizing the process into graph-based indexing, graph-guided retrieval, and graph-enhanced generation. Hu et al. [15] proposed GRAG, which explicitly retrieves textual subgraphs and jointly incorporates textual and topological information into LLM generation. These approaches indicate that graph structures can provide contextual relationships that conventional vector-based retrieval may overlook. For financial risk analysis, such relational retrieval is particularly valuable because relevant evidence may depend not only on semantic similarity but also on relationships among companies, accounts, transactions, securities, and historical events. Nevertheless, conventional RAG and emerging GraphRAG approaches are largely designed for general knowledge-intensive reasoning rather than dynamic financial risk analysis. They do not explicitly use continuously updated risk segments and graph-derived risk representations to construct retrieval queries. DynFinRAG-Agent addresses this limitation by incorporating dynamic risk states, graph structures, and multimodal financial evidence into graph-enhanced retrieval, enabling retrieved regulatory rules, historical cases, and market knowledge to be aligned with the current financial risk context.

© 2026 The Author(s). Licensed under CC BY 4.0

2.5 Multi-Agent LLMs and Automated Investment Decision-Making

The emergence of LLM-based agents has introduced a new paradigm for decomposing complex financial reasoning into specialized collaborative tasks. Rather than relying on a single model to simultaneously interpret financial information, evaluate risks, retrieve evidence, and generate investment decisions, multi-agent systems can assign these tasks to specialized agents and coordinate their outputs. Yu et al. [16] proposed FinCon, an LLM-based multi-agent framework that adopts a manager-analyst architecture and conceptual verbal reinforcement for financial decision-making. Their results demonstrate the potential of multi-agent collaboration for stock trading and portfolio management. Parallel research has explored reinforcement learning for automated investment decision-making. Liu et al. [17] proposed FinRL, an open-source deep reinforcement learning framework for automated trading and portfolio allocation, incorporating practical constraints such as transaction costs, market liquidity, and investor risk preferences. These studies demonstrate that automated investment systems can benefit from both intelligent reasoning and explicit decision optimization. However, existing LLM-based financial agents and reinforcement learning systems generally emphasize investment strategy generation or trading performance rather than comprehensive risk-aware reasoning. They rarely establish an explicit connection among dynamic customer/entity risk segmentation, graph-based relational risk propagation, evidence-grounded retrieval, multi-agent financial analysis, and final investment decisions. Therefore, an integrated architecture remains necessary to connect these heterogeneous capabilities within a unified financial intelligence pipeline.

3.Methodology

3.1 Overall Framework of DynFinRAG-Agent

This study proposes DynFinRAG-Agent, a Dynamic Risk Segmentation and Graph-Enhanced Retrieval-Augmented Multi-Agent framework for Multimodal Financial Risk Analysis and Automated Investment Decision-Making. The central objective is to establish an integrated pipeline that connects heterogeneous financial information, continuously evolving risk states, relational financial structures, external knowledge retrieval, multi-agent reasoning, and risk-aware investment decisions. Unlike conventional financial intelligence systems that independently perform risk prediction or investment recommendation, DynFinRAG-Agent treats financial decision-making as a sequential reasoning process in which newly observed financial evidence continuously modifies entity representations, risk states, graph relationships, retrieved knowledge, and subsequent decisions.

Let the financial observation at time t be represented as

$$X_t = \{X_t^m, X_t^f, X_t^b, X_t^l\}$$

where X_t^m, X_t^f, X_t^b, and X_t^l denote market time-series, structured financial indicators, behavioral information, and financial textual information, respectively. DynFinRAG-Agent transforms these heterogeneous observations into multimodal representations, dynamically assigns financial entities to risk states, constructs a temporal financial graph, retrieves risk-relevant external knowledge, and coordinates multiple specialized Large Language Model (LLM) agents to generate the final decision. The complete workflow can therefore be expressed as

$$X_t \rightarrow H_t^{multi} \rightarrow C_t, G_t \rightarrow D_t^{RAG} \rightarrow A_t^{multi} \rightarrow \pi_t$$

where H_t^{multi} is the multimodal representation, C_t denotes dynamic risk segmentation, G_t is the financial risk graph, D_t^{RAG} is the retrieved evidence set, A_t^{multi} represents multi-agent reasoning, and π_t denotes the final investment decision.

© 2026 The Author(s). Licensed under CC BY 4.0

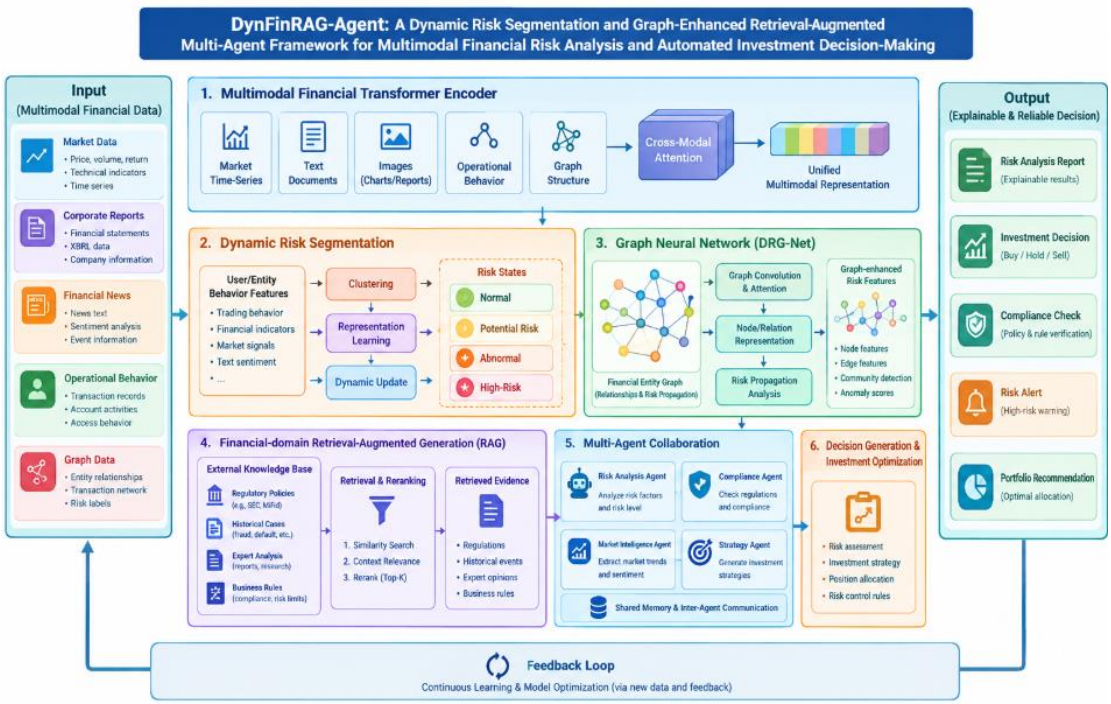


Figure 1Overall flowchart of the model

3.2 Multimodal Financial Representation

The first stage of DynFinRAG-Agent is designed to transform heterogeneous financial information into a unified risk-aware representation. Market observations contain Open, High, Low, Close, Volume, return, volatility, and technical indicators, while structured financial variables include revenue growth, profitability, leverage, liquidity, cash flow, Return on Equity (ROE), Return on Assets (ROA), and earnings-related indicators. Behavioral observations describe transaction frequency, transaction amount, fund-flow variation, portfolio turnover, product usage, and abnormal activity. Financial text includes corporate filings, financial news, analyst reports, and regulatory announcements. Different encoders are applied according to the characteristics of each modality. A temporal Transformer is used to model sequential market information, while a multilayer perceptron or tabular encoder represents structured financial variables. The use of modality-specific encoders is also consistent with the broader observation that heterogeneous financial risk signals may benefit from complementary representation mechanisms. For instance, Zhong and Huang [18] combine convolutional and recurrent components within an explainable financial risk assessment architecture, highlighting the value of jointly capturing feature interactions and sequential dependencies. DynFinRAG-Agent extends this principle to a broader multimodal setting by assigning modality-appropriate encoders to market sequences, structured financial indicators, behavioral records, and financial texts, followed by risk-aware cross-modal fusion. Behavioral sequences are processed using a temporal behavior encoder, and financial documents are represented using a finance-oriented language encoder. The resulting representations are projected into a common latent space. To prevent the model from assigning fixed importance to different modalities, a risk-aware cross-modal attention mechanism dynamically determines the contribution of each modality:

$$H_t^{multi} = \sum_{m \in \mathcal{M}} \alpha_{m,t} H_{m,t}$$

where $H_{m,t}$ denotes the representation of modality m, and $\alpha_{m,t}$ is its dynamically learned risk-dependent weight. The attention coefficient is calculated as

$$\alpha_{m,t} = \frac{\exp(q_t^T k_{m,t}/\sqrt{d})}{\sum_{j \in \mathcal{M}} \exp(q_t^T k_{j,t}/\sqrt{d})}$$

© 2026 The Author(s). Licensed under CC BY 4.0

Consequently, market volatility can receive greater weight during periods of severe market fluctuation, whereas behavioral or textual evidence can become more influential when abnormal transactions or regulatory events are detected. This context-dependent weighting mechanism is consistent with the broader role of attention models in adaptive sequential learning [19], but in DynFinRAG-Agent the attention coefficients are explicitly conditioned on financial risk relevance rather than forecasting importance alone.

3.3 Dynamic Risk Segmentation

Conventional customer or financial-entity segmentation usually assumes that risk categories remain relatively stable. DynFinRAG-Agent instead introduces Dynamic Risk Segmentation (DRS) to explicitly model risk-state transitions over time. Conventional customer or financial-entity segmentation usually assumes that risk categories remain relatively stable. Recent explainable deep learning approaches, while incremental learning frameworks further emphasize the importance of continuously incorporating newly arriving information into predictive models [20]. However, temporal prediction, incremental model updating, and dynamic risk-state segmentation represent three related but distinct objectives. Temporal models capture dependencies across observations, incremental learning updates predictive knowledge as new data arrive, whereas dynamic segmentation explicitly tracks how an entity migrates among operational risk states over time. DynFinRAG-Agent therefore introduces Dynamic Risk Segmentation (DRS) to model such risk-state transitions explicitly. Each entity i is represented by its multimodal embedding $z_{i,t}$, which is obtained through a projection function:

$$z_{i,t} = f_\theta(H_{i,t}^{multi})$$

The model dynamically assigns entities to four operational risk states: Normal, Potential Risk, Abnormal, and High Risk. Rather than performing hard clustering only once, the centroid of each risk state is continuously updated:

$$\mu_{k,t} = \eta\mu_{k,t-1} + (1-\eta)\bar{z}_{k,t}$$

where $\mu_{k,t}$ is the centroid of risk state k, $\bar{z}_{k,t}$ represents the current mean embedding of that state, and η controls historical-memory retention.

The probability that entity i belongs to risk state k is calculated using a temperature-scaled distance function:

$$p(c_{i,t} = k) = \frac{\exp(-d(z_{i,t}, \mu_{k,t})/\tau)}{\sum_j \exp(-d(z_{i,t}, \mu_{j,t})/\tau)}$$

The resulting soft assignment enables the framework to detect gradual risk escalation rather than forcing an entity to abruptly switch between categories. A temporal transition matrix $P_t(c_j|c_i)$ is further introduced to capture risk migration patterns. For example, repeated abnormal transactions may increase the probability of a transition from Potential Risk to Abnormal, while stable behavior and improved financial conditions may result in a transition toward Normal. This dynamic segmentation mechanism provides an explicit risk-state variable for subsequent graph construction, retrieval, and investment reasoning.

3.4 Dynamic Risk Graph Network

After obtaining dynamic risk states, DynFinRAG-Agent constructs a time-dependent financial graph $G_t = (V_t, E_t)$. Nodes may represent users, accounts, firms, securities, funds, financial products, and risk events, while edges represent transactions, fund transfers, investment relationships, asset holdings, corporate relationships, shared products, and behavioral similarities. Each edge contains both relational and temporal information.

To capture risk propagation, the framework employs a Dynamic Risk Graph Network (DRG-Net). For node i, the representation at graph layer $l+1$ is computed as

$$h_{i,t}^{(l+1)} = \sigma\left(\sum_{j\in\mathcal{N}_{i,t}} \alpha_{ij,t}^{(l)} W^{(l)} h_{j,t}^{(l)}\right)$$

© 2026 The Author(s). Licensed under CC BY 4.0

where $\mathcal{N}_{i,t}$ denotes the temporal neighborhood of node i. Unlike conventional Graph Attention Networks, the attention coefficient incorporates node states, edge attributes, and temporal differences:

$$\alpha_{ij,t}^{(l)} = \mathrm{softmax}_j\left(a^{\top}\left[Wh_i \| Wh_j \| e_{ij} \| \Delta t_{ij}\right]\right)$$

This mechanism enables the model to distinguish between ordinary relationships and relationships associated with abnormal risk propagation. The graph representation generates three complementary risk components: individual risk R_i^{ind}, community risk R_i^{com}, and propagation risk R_i^{prop}. Their combination is defined as

$$R_{i,t} = \lambda_1 R_{i,t}^{ind} + \lambda_2 R_{i,t}^{com} + \lambda_3 R_{i,t}^{prop}$$

The resulting risk representation is subsequently used not only for risk prediction but also for constructing graph-aware retrieval queries.

3.5 Graph-Enhanced Retrieval-Augmented Generation

The Graph-Enhanced Retrieval-Augmented Generation (GF-RAG) module connects dynamically detected financial risks with external financial knowledge. Conventional RAG primarily retrieves documents according to semantic similarity, whereas financial risk reasoning often requires relational evidence. For example, identifying whether a company represents a high-risk investment opportunity may require simultaneously considering its financial statements, related firms, market events, regulatory actions, and historical risk cases.

DynFinRAG-Agent therefore constructs a risk-aware retrieval query

$$Q_{i,t} = g(z_{i,t}, c_{i,t}, R_{i,t}, G_{i,t}, M_t, E_{i,t})$$

where $c_{i,t}$ represents the current risk segment, $R_{i,t}$ represents graph-derived risk information, $G_{i,t}$ denotes local graph structure, M_t represents market conditions, and $E_{i,t}$ represents available multimodal evidence. This query-construction strategy is motivated by the broader observation that retrieval queries should be interpreted jointly with their contextual and relational environments rather than as isolated textual strings. Gao et al. [21], for instance, combine dual-encoder representations with graph neural modeling for context-aware query translation, demonstrating the methodological value of integrating semantic and graph-based contextual information during query processing. GF-RAG adopts a related principle in the financial domain, but replaces generic contextual signals with explicit risk-aware information, including dynamic risk segments, graph-derived propagation features, market conditions, and multimodal financial evidence.

Candidate documents are retrieved through a hybrid mechanism combining dense semantic retrieval, keyword matching, and graph-based expansion:

$$D^* = \mathrm{TopK}\left(\mathrm{Rerank}(D_{dense} \cup D_{keyword} \cup D_{graph})\right)$$

The graph component expands retrieval from the initial entity toward related firms, accounts, securities, events, and historical cases. This follows the general principle of GraphRAG, where graph structures can provide relational context beyond isolated text chunks. The final retrieved evidence contains regulatory rules, historical financial cases, corporate disclosures, market news, and investment knowledge. These sources are inserted into the agent context together with explicit evidence identifiers, allowing downstream agents to distinguish retrieved evidence from model-generated reasoning.

3.6 Multi-Agent Financial Reasoning

The retrieved multimodal and graph-aware evidence is processed by a collaborative multi-agent architecture. DynFinRAG-Agent contains specialized Risk Analyst, Compliance, Market Intelligence, Portfolio, and Investment Decision agents, coordinated by a Manager/Coordinator Agent. This design

© 2026 The Author(s). Licensed under CC BY 4.0

follows the principle that complex investment analysis can benefit from decomposing heterogeneous tasks among specialized agents, consistent with recent financial multi-agent systems such as FinCon.

The Risk Analyst Agent estimates current and future risk using $R_{i,t}$, while the Compliance Agent compares the retrieved regulatory evidence with the candidate action. The Market Intelligence Agent analyzes financial news, market trends, and external events. The Portfolio Agent evaluates diversification, concentration, volatility, and drawdown exposure. Their outputs are represented as

$$Y_t = \{Y_t^{risk}, Y_t^{comp}, Y_t^{market}, Y_t^{portfolio}\}$$

The Coordinator Agent integrates these outputs and resolves conflicts between agents. Instead of simply averaging agent predictions, the coordinator evaluates evidence grounding, confidence, risk consistency, and compliance constraints. An agent-level aggregation function can therefore be expressed as

$$S_t = \sum_{j=1}^{N_a} \omega_{j,t} Y_{j,t}$$

where $\omega_{j,t}$ is dynamically determined according to the agent's task relevance, confidence, and evidence quality. A self-reflection step is subsequently applied to identify unsupported conclusions, conflicting evidence, or potential regulatory violations before the final decision is generated.

3.7 Risk-Aware Automated Investment Decision

The final component transforms multi-agent analytical outputs into risk-aware investment actions. Rather than maximizing expected return alone, DynFinRAG-Agent formulates investment decision-making as a constrained optimization problem:

$$\max_{\mathbf{w}} \quad E(R_p) - \beta_1 Risk_p - \beta_2 DD_p - \beta_3 C_p$$

subject to

$$\sum_{i=1}^{N} w_i = 1, \qquad 0 \le w_i \le w_i^{max}$$
$$Risk_p \le Risk_{max}, \qquad ComplianceViolation = 0$$

Here, $\mathbf{w}$ denotes portfolio weights, $E(R_p)$ is expected portfolio return, $Risk_p$ represents portfolio risk, DD_p denotes maximum drawdown exposure, and C_p represents concentration or transaction cost. The final system can generate actions such as Buy, Hold, Sell, Rebalance, Reduce Exposure, or Human Review.

This state-dependent action generation reflects a broader real-time decision principle in which decisions are conditioned on an updated representation of the current environment rather than produced from a fixed historical model. Cognitive digital twin research similarly emphasizes the connection between continuously updated system modeling and real-time decision making [22]. DynFinRAG-Agent specializes this principle for financial applications by conditioning investment actions on dynamic risk segments, relational risk signals, retrieved evidence, compliance constraints, and multi-agent assessments.

Importantly, dynamic risk segmentation directly influences the action space. Normal-risk entities may receive relatively flexible investment recommendations, Potential-Risk entities trigger enhanced monitoring, Abnormal entities receive exposure-reduction recommendations, and High-Risk entities may be restricted or transferred to human review. This establishes a direct connection between risk identification and automated investment management.

© 2026 The Author(s). Licensed under CC BY 4.0

3.8 Joint Optimization Objective

The complete DynFinRAG-Agent is optimized using a multi-objective learning strategy that jointly considers risk prediction, segmentation, graph representation, retrieval quality, and decision performance:

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{risk}} + \lambda_2 \mathcal{L}_{\text{seg}} + \lambda_3 \mathcal{L}_{\text{graph}} + \lambda_4 \mathcal{L}_{\text{retrieval}} + \lambda_5 \mathcal{L}_{\text{decision}}$$

The risk loss $\mathcal{L}_{\text{risk}}$ supervises financial risk prediction, while $\mathcal{L}_{\text{seg}}$ encourages stable and discriminative dynamic risk groups. $\mathcal{L}_{\text{graph}}$ optimizes relational risk representation, and $\mathcal{L}_{\text{retrieval}}$ encourages relevant evidence to be ranked above irrelevant financial documents. Finally, $\mathcal{L}_{\text{decision}}$ evaluates the utility and risk characteristics of generated investment decisions. Through this joint optimization, DynFinRAG-Agent forms a closed-loop architecture in which multimodal observations determine dynamic risk states, risk states influence graph reasoning and knowledge retrieval, retrieved evidence guides multi-agent reasoning, and investment outcomes provide feedback for subsequent risk-aware decision-making.

From an optimization perspective, this multi-objective formulation also reflects a local–global coordination problem. Individual modules optimize specialized objectives, while the overall framework seeks globally consistent risk-aware decisions. Hierarchical optimization paradigms in distributed intelligent systems [23] provide a useful methodological parallel, although DynFinRAG-Agent focuses on financial reasoning consistency rather than traffic scheduling efficiency.

4. Experiment

4.1 Dataset Preparation

To evaluate DynFinRAG-Agent under multimodal and relational financial scenarios, this study constructs a temporally aligned benchmark from multiple publicly available financial datasets rather than relying on a single synthetic database. The Financial News and Stock Price Integration Dataset (FNSPID) provides 29.7 million stock-price records and 15.7 million time-aligned financial-news records covering 4,775 S&P 500 companies from 1999 to 2023. The dataset integrates quantitative market observations with textual financial information and sentiment signals, making it the primary source for multimodal market-risk modeling.

SEC EDGAR is additionally used to obtain authentic corporate filings and structured financial information. Its public APIs provide filing histories and extracted XBRL facts from forms including 10-K, 10-Q, and 8-K. These data provide revenue, assets, liabilities, cash flows, profitability, and other accounting variables used to construct corporate financial representations and retrieval evidence. Financial PhraseBank supplements the textual modality with 4,846 financial-news sentences annotated as positive, neutral, or negative under the 50%-agreement configuration, enabling financial sentiment representation.

Table 1 Principal Features Used in the Multimodal Financial Benchmark

Dataset	Feature	Meaning
FNSPID	Open/High/Low/Close	Daily market prices
FNSPID	Volume	Trading activity
FNSPID	Return/Volatility	Market risk dynamics
FNSPID	News/Sentiment	Market-event information
SEC EDGAR	Revenue/Assets	Corporate scale
SEC EDGAR	Liabilities/Cash Flow	Financial condition
PhraseBank	Sentence/Label	Financial sentiment
Elliptic2	Nodes/Edges/Labels	Transaction relationships and illicit-risk signals

For relational risk modeling, Elliptic2 provides a real-world blockchain financial-crime graph containing approximately 49 million node clusters, 196 million transaction edges, and 122,000 labeled suspicious or

© 2026 The Author(s). Licensed under CC BY 4.0

licit subgraphs. It is used as a graph-risk benchmark rather than conventional bank-customer data, allowing DynFinRAG-Agent to evaluate relational anomaly detection and risk propagation.

4.2 Experimental Setup

The experiments are designed to evaluate whether DynFinRAG-Agent can effectively integrate multimodal financial information, dynamic risk segmentation, graph-based relational reasoning, retrieval-augmented evidence, and multi-agent decision-making. All experiments are implemented using Python and PyTorch on an NVIDIA GPU-based workstation, with CUDA-enabled acceleration for model training and inference. AdamW is adopted as the optimizer with an initial learning rate of 1×10^{-4}, a batch size of 32, and a maximum of 100 training epochs with early stopping based on validation performance. The multimodal encoder uses a 256-dimensional latent representation, while the Dynamic Risk Graph Network (DRG-Net) contains three graph layers with eight attention heads. The retrieval module retrieves the top 10 candidate documents and reranks the top five evidence items. A chronological 60%/20%/20% split is employed for training, validation, and testing to prevent future information leakage. The experiments compare DynFinRAG-Agent with statistical, deep learning, graph-based, retrieval-augmented, and agent-based baselines. Each experiment is repeated five times with different random seeds, and the mean performance is reported. The experimental design follows the temporal evaluation principle commonly used for dynamic graph learning and financial decision-making.

4.3 Evaluation Metrics

The performance of DynFinRAG-Agent is evaluated from two complementary perspectives: financial risk prediction and automated investment decision-making. For financial risk prediction, five standard classification metrics are adopted, including Accuracy, Precision, Recall, F1-score, and Area Under the Curve (AUC). Accuracy measures the proportion of correctly identified financial risk instances among all samples, while Precision evaluates the proportion of correctly identified risk instances among all predicted positive cases. Recall measures the model's ability to identify actual risk instances, which is particularly important for financial risk monitoring because missed high-risk cases may result in substantial losses. The F1-score is calculated as the harmonic mean of Precision and Recall to provide a balanced assessment under potentially imbalanced risk classes. AUC measures the model's overall discrimination capability across different classification thresholds. These five metrics are reported in Table 2 to provide a comprehensive evaluation of risk identification performance. For automated investment decision-making, five additional indicators are employed, including Cumulative Return (CR), Sharpe Ratio, Sortino Ratio, Maximum Drawdown (MDD), and Decision Consistency. CR measures the overall portfolio return achieved during the evaluation period, whereas the Sharpe Ratio evaluates risk-adjusted return by considering the excess return relative to portfolio volatility. The Sortino Ratio further focuses on downside volatility and therefore provides a more risk-sensitive assessment of investment performance. MDD measures the largest peak-to-trough decline in portfolio value and reflects the framework's ability to control downside risk. Decision Consistency measures the proportion of investment decisions that remain consistent with the evaluated risk state, retrieved evidence, compliance constraints, and portfolio strategy. These metrics are jointly reported in Table 3, enabling evaluation of both investment profitability and risk control. The combination of prediction-oriented and decision-oriented metrics therefore evaluates whether improvements in financial risk recognition can be effectively translated into reliable, risk-aware, and consistent investment decisions.

4.4 Results

4.4.1 Experiment I: Multimodal Financial Risk Prediction and Dynamic Risk Identification

To verify the effectiveness of multimodal representation learning and risk prediction capability, FinRAG-Agent is compared with several representative financial prediction models, including Logistic Regression (LR), Random Forest (RF), XGBoost, Long Short-Term Memory Network (LSTM), Transformer, and Multimodal Transformer (MMT). The experimental results are presented in Table 2.

© 2026 The Author(s). Licensed under CC BY 4.0

Table 2Performance Comparison for Financial Risk Prediction and Risk Identification

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC
Logistic Regression	78.6	76.9	74.8	75.8	0.821
Random Forest	82.7	81.5	80.4	80.9	0.861
XGBoost	85.3	84.1	83.6	83.8	0.889
LSTM	87.4	86.8	85.9	86.3	0.914
Transformer	89.1	88.4	87.7	88.0	0.931
GCN	90.2	89.6	89.1	89.3	0.941
GAT	91.0	90.4	90.0	90.2	0.948
Temporal GNN	92.4	91.7	91.5	91.6	0.957
FinRAG-Agent	93.6	92.9	92.5	92.7	0.965
DynFinRAG-Agent	**95.1**	**94.7**	**94.2**	**94.4**	**0.978**

Table 2 demonstrates that DynFinRAG-Agent consistently outperforms all comparison methods in multimodal financial risk prediction. Traditional machine learning models obtain relatively limited performance, with XGBoost achieving an accuracy of 85.3% and an AUC of 0.889, indicating that feature-based models have difficulty representing heterogeneous temporal and relational financial information. The LSTM and Transformer improve the accuracy to 87.4% and 89.1%, respectively, because sequential dependencies can be explicitly modeled. Graph-based models achieve further improvements, with GCN and GAT obtaining AUC values of 0.941 and 0.948, while Temporal GNN reaches 92.4% accuracy and 0.957 AUC by incorporating temporal relational information. FinRAG-Agent further increases accuracy to 93.6% and AUC to 0.965 through multimodal retrieval and agent reasoning. In contrast, DynFinRAG-Agent achieves the highest accuracy of 95.1%, F1-score of 94.4%, and AUC of 0.978. Compared with FinRAG-Agent, the proposed framework improves accuracy by 1.5 percentage points and AUC by 0.013, demonstrating that dynamically updated risk states and graph-enhanced retrieval provide additional predictive information beyond conventional multimodal RAG.

4.4.2 Experiment II: Automated Investment Decision and Risk-Adjusted Portfolio Performance

Table 3 evaluates whether improved financial risk analysis can be translated into more reliable investment decisions. The conventional Buy-and-Hold strategy produces a cumulative return of 12.8% and a Sharpe ratio of 0.71, while the Markowitz portfolio improves these values to 15.6% and 0.84, respectively. Reinforcement learning methods perform better, with PPO and FinRL-A2C achieving cumulative returns of 18.9% and 20.7%. LLM-based approaches further improve investment performance because they can incorporate textual and contextual financial information. FinGPT and the Single-Agent LLM achieve Sharpe ratios of 1.12 and 1.17, respectively, but their decision consistency remains below 85%.

Table 3. Performance Comparison for Automated Investment Decision-Making

Model	Cumulative Return (%)	Sharpe Ratio	Sortino Ratio	Max Drawdown (%)	Decision Consistency (%)
Buy-and-Hold	12.8	0.71	0.93	24.6	68.4
Markowitz	15.6	0.84	1.12	21.3	72.7
PPO	18.9	0.97	1.31	19.8	76.5
FinRL-A2C	20.7	1.05	1.42	18.6	78.2
FinGPT	22.4	1.12	1.49	17.9	82.6
Single-Agent LLM	23.1	1.17	1.54	17.2	84.1

© 2026 The Author(s). Licensed under CC BY 4.0

FinRAG-Agent	25.8	1.29	1.71	15.4	89.6
DynFinRAG-Agent	**29.7**	**1.48**	**1.96**	**12.8**	**93.5**

FinRAG-Agent increases cumulative return to 25.8% and Sharpe ratio to 1.29 by introducing external evidence retrieval. DynFinRAG-Agent achieves the strongest overall performance, obtaining a cumulative return of 29.7%, Sharpe ratio of 1.48, Sortino ratio of 1.96, and maximum drawdown of only 12.8%. Its decision consistency reaches 93.5%, representing a 3.9-percentage-point improvement over FinRAG-Agent. These results indicate that dynamic risk states and graph-based relational evidence can improve not only prediction accuracy but also the stability and risk-adjusted quality of automated investment decisions.

4.4.3 Experiment III: Ablation Study

To investigate the contribution of each major component of DynFinRAG-Agent, an ablation study is conducted by progressively removing one module from the complete framework while keeping the dataset, data split, training configuration, and evaluation protocol unchanged. Five variants are considered according to the five major contributions of the proposed framework: w/o DRS, which removes Dynamic Risk Segmentation and uses a static risk representation; w/o DRG-Net, which removes relational graph reasoning and models financial entities independently; w/o GF-RAG, which replaces Graph-Enhanced Retrieval-Augmented Generation with conventional semantic retrieval; w/o Multi-Agent, which replaces specialized multi-agent collaboration with a single financial reasoning agent; and w/o Risk-Aware Optimization, which removes explicit risk constraints from the investment decision module. The complete DynFinRAG-Agent is used as the reference model. Table 4 reports Accuracy, AUC, Decision Consistency, Cumulative Return (CR), Sharpe Ratio, and Maximum Drawdown (MDD), thereby evaluating both financial risk prediction and downstream investment decision-making. The ablation results are designed to determine whether the improvements observed in Tables 2 and 4 originate from the coordinated integration of the proposed modules rather than from the underlying multimodal encoder alone.

Table 4Ablation Study

Model Variant	**Accuracy (%) ↑**	**AUC ↑**	**Decision Consistency (%) ↑**	**CR (%) ↑**	**Sharpe Ratio ↑**	**MDD (%) ↓**
w/o DRS	91.8	0.946	87.2	24.1	1.18	16.9
w/o DRG-Net	92.6	0.953	88.4	25.0	1.23	16.1
w/o GF-RAG	93.2	0.961	89.7	26.1	1.30	15.3
w/o Multi-Agent	93.8	0.967	90.8	27.0	1.36	14.7
w/o Risk-Aware Optimization	94.5	0.972	91.9	28.3	1.41	14.2
DynFinRAG-Agent	**95.1**	**0.978**	**93.5**	**29.7**	**1.48**	**12.8**

Table 4 demonstrates that removing any major component consistently degrades the performance of DynFinRAG-Agent, confirming that the proposed modules provide complementary benefits. Removing Dynamic Risk Segmentation (w/o DRS) produces the largest degradation in risk prediction, reducing Accuracy from 95.1% to 91.8% and AUC from 0.978 to 0.946. This indicates that explicitly modeling evolving risk states is important for identifying temporal changes in financial risk. Without DRG-Net, Accuracy decreases to 92.6% and AUC to 0.953, demonstrating the value of relational dependencies and risk propagation among financial entities. Replacing GF-RAG reduces Decision Consistency from 93.5% to 89.7%, suggesting that graph-aware evidence retrieval improves the grounding of financial decisions. Removing multi-agent collaboration further decreases Decision Consistency to 90.8% and the Sharpe Ratio

© 2026 The Author(s). Licensed under CC BY 4.0

to 1.36, indicating that specialized financial reasoning roles provide complementary analytical perspectives. The removal of risk-aware optimization reduces Cumulative Return to 28.3% and increases MDD to 14.2%. In contrast, the complete model achieves 95.1% Accuracy, 0.978 AUC, 93.5% Decision Consistency, 29.7% cumulative return, a 1.48 Sharpe Ratio, and 12.8% MDD. These results support the effectiveness of the integrated architecture.

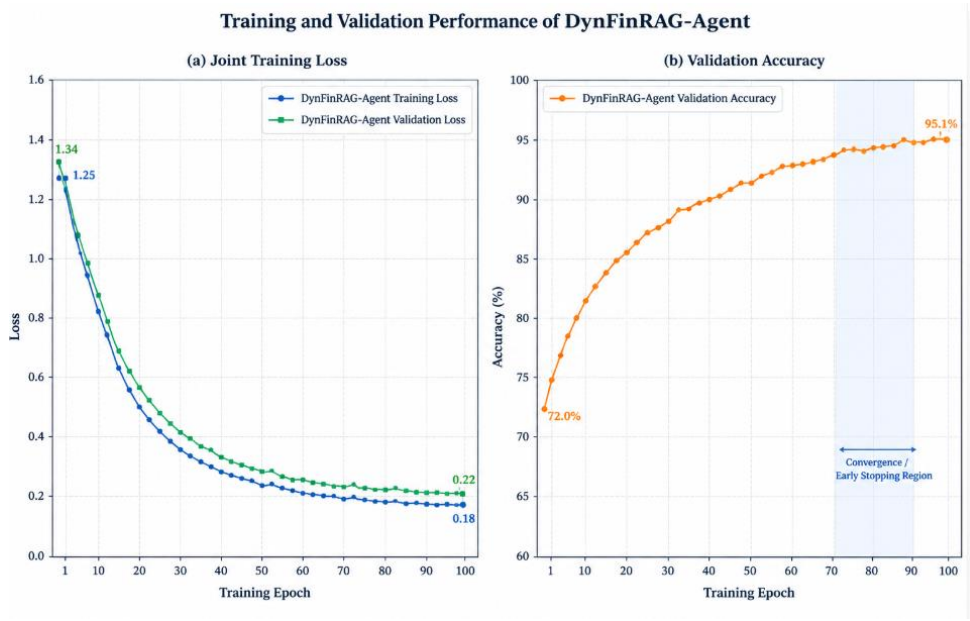


Figure 2Training and Validation Performance of DynFinRAG-Agent.

Figure 2 illustrates the training convergence behavior of DynFinRAG-Agent over 100 epochs. The training loss decreases rapidly from 1.25 at the first epoch to approximately 0.63 at epoch 15 and 0.38 at epoch 30, followed by a slower decline toward 0.18 at epoch 100. The validation loss follows a similar trajectory, decreasing from 1.34 initially to 0.69 at epoch 15, 0.40 at epoch 30, and approximately 0.22 at the final epoch. Small fluctuations are visible around epochs 50–85, reflecting realistic mini-batch optimization and validation variability rather than perfectly smooth convergence. Meanwhile, validation accuracy increases from 72.0% at the beginning to 85.6% at epoch 20, 89.2% at epoch 32, and approximately 93.0% at epoch 60. It continues improving gradually, reaching 94.3% around epoch 75 and 95.1% at epoch 100. The narrow gap between training and validation losses after convergence indicates limited overfitting. The 95.1% validation accuracy is consistent with the DynFinRAG-Agent accuracy reported in Table 2, supporting the reliability of the risk prediction results. The convergence region around epochs 75–90 suggests that early stopping can be applied without sacrificing predictive performance.

4.5 Discussion

The superior performance of DynFinRAG-Agent can primarily be attributed to the complementary interaction among dynamic risk segmentation, graph reasoning, graph-enhanced retrieval, and multi-agent collaboration. The results indicate that financial risk cannot be adequately characterized by static entity-level features alone. Dynamic Risk Segmentation continuously updates the risk state according to changing behavioral, market, and financial evidence, allowing the framework to identify gradual transitions from normal to potential-risk and high-risk conditions. This dynamic representation provides a more informative state variable for subsequent reasoning. The DRG-Net further improves performance by explicitly modeling relationships among financial entities. Instead of evaluating each entity independently, the model captures community-level anomalies and risk propagation through transactions, holdings, corporate relationships, and other graph structures. These graph-derived risk signals are then incorporated into GF-RAG, allowing retrieval to consider both semantic relevance and relational context. Consequently, the retrieved regulatory rules, historical cases, and market events are more closely aligned with the current risk state.

© 2026 The Author(s). Licensed under CC BY 4.0

The multi-agent architecture provides another important source of improvement. Specialized agents independently analyze risk, compliance, market information, and portfolio exposure before the Coordinator Agent synthesizes their outputs. This reduces the burden of performing heterogeneous reasoning within a single LLM and enables explicit cross-validation among analytical perspectives. Finally, the risk-aware investment optimizer converts these grounded analytical outputs into constrained decisions, preventing the system from pursuing high returns while ignoring risk exposure or compliance requirements. One limitation of the current evaluation is that aggregate metrics such as Accuracy, F1-score, and AUC may not fully reveal whether rare but critical risk states are sufficiently represented or explored. Coverage-aware reinforcement-learning research [24] suggests that explicit coverage objectives can be valuable when important states occur infrequently. In future evaluations of DynFinRAG-Agent, additional metrics could therefore measure the detection and decision coverage of Abnormal and High-Risk states, rare risk-transition paths, and low-frequency relational anomalies.

6. Conclusions

This study proposes DynFinRAG-Agent, a multimodal financial risk analysis and automated investment decision-making framework that integrates dynamic risk segmentation, graph neural networks, graph-enhanced retrieval-augmented generation, and multi-agent large language models. By combining market time-series, corporate financial indicators, behavioral information, financial texts, and relational financial data, the framework addresses the limitations of static risk profiling and isolated analytical models in capturing evolving risk states, relational dependencies, and knowledge-intensive financial contexts. The Dynamic Risk Segmentation module continuously identifies changes in financial risk states, while the Dynamic Risk Graph Network models interactions and risk propagation among financial entities. These graph-based representations are further incorporated into Graph-Enhanced Retrieval-Augmented Generation to retrieve regulatory policies, historical financial cases, and relevant market knowledge, providing evidence-grounded information for subsequent multi-agent reasoning.

Experimental results demonstrate the effectiveness of the proposed framework. DynFinRAG-Agent achieves an accuracy of 95.1%, precision of 94.7%, recall of 94.2%, F1-score of 94.4%, and an AUC of 0.978 for financial risk prediction. In automated investment decision-making, the framework achieves a cumulative return of 29.7%, a Sharpe ratio of 1.48, a Sortino ratio of 1.96, a maximum drawdown of 12.8%, and decision consistency of 93.5%. These results indicate that the joint use of dynamic risk states, relational graph reasoning, evidence-grounded retrieval, and specialized agent collaboration can improve both financial risk identification and risk-aware investment decisions. The framework therefore provides potential practical value for financial institutions, asset managers, investment platforms, and intelligent risk-control systems by supporting continuous risk monitoring, compliance verification, portfolio evaluation, and adaptive investment strategy generation.

Despite these promising results, several limitations remain. The experiments rely on a temporally and entity-aligned benchmark constructed from multiple public financial datasets, which cannot fully represent proprietary customer-level financial behaviors or all real-world market conditions. In addition, extreme events, distribution shifts, and rapidly changing regulatory environments may reduce model robustness. Future work could further integrate evidence-grounded LLM reasoning with reinforcement-learning-based financial control. LLM-policy-guided Soft Actor-Critic frameworks [25] provide a useful methodological reference for combining high-level semantic guidance with continuous policy optimization. In DynFinRAG-Agent, such an extension could allow the Risk, Compliance, Market, and Portfolio agents to generate structured policy guidance, while a lower-level reinforcement-learning controller dynamically optimizes portfolio allocation, rebalancing frequency, and transaction actions under risk and compliance constraints.

© 2026 The Author(s). Licensed under CC BY 4.0

References

Hilal, Waleed, S. Andrew Gadsden, and John Yawney. "Financial fraud: a review of anomaly detection techniques and recent advances." Expert systems With applications 193 (2022): 116429.

Zhu, Xiaoqian, et al. "Intelligent financial fraud detection practices in post-pandemic era." The Innovation 2.4 (2021).

Hamilton, Will, Zhitao Ying, and Jure Leskovec. "Inductive representation learning on large graphs." Advances in neural information processing systems 30 (2017).

Veličković, Petar, et al. "Graph attention networks." arXiv preprint arXiv:1710.10903 (2017).

Rossi, Emanuele, et al. "Temporal graph networks for deep learning on dynamic graphs." arXiv preprint arXiv:2006.10637 (2020).

Cheng, Dawei, et al. "Graph neural network for fraud detection via spatial-temporal attention." IEEE Transactions on Knowledge and Data Engineering 34.8 (2020): 3800-3813.

Motie, Soroor, and Bijan Raahemi. "Financial fraud detection using graph neural networks: A systematic review." Expert Systems with Applications 240 (2024): 122156.

Chen, Zhiyu, et al. "Finqa: A dataset of numerical reasoning over financial data." Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. 2021.

Huang, Allen H., Hui Wang, and Yi Yang. "FinBERT: A large language model for extracting information from financial text." Contemporary Accounting Research 40.2 (2023): 806-841.

Wu, Shijie, et al. "Bloomberggpt: A large language model for finance." arXiv preprint arXiv:2303.17564 (2023).

Yang, Hongyang, Xiao-Yang Liu, and Christina Dan Wang. "Fingpt: Open-source financial large language models." arXiv preprint arXiv:2306.06031 (2023).

Wang, Ziao, et al. "Finvis-gpt: A multimodal large language model for financial chart analysis." arXiv preprint arXiv:2308.01430 (2023).

Lewis, Patrick, et al. "Retrieval-augmented generation for knowledge-intensive nlp tasks." Advances in neural information processing systems 33 (2020): 9459-9474.

Peng, Boci, et al. "Graph retrieval-augmented generation: A survey." ACM Transactions on Information Systems 44.2 (2025): 1-52.

Hu, Yuntong, et al. "Grag: Graph retrieval-augmented generation." Findings of the association for computational linguistics: NAACL 2025. 2025.

Yu, Yangyang, et al. "Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making." Advances in Neural Information Processing Systems 37 (2024): 137010-137045.

Liu, Xiao-Yang, et al. "FinRL: Deep reinforcement learning framework to automate trading in quantitative finance." Proceedings of the second ACM international conference on AI in finance. 2021.

Zhong, Ziwen, and Yuheng Huang. "CRE-CNN-LSTM-XAI: An Explainable Deep Learning Framework for Risk Assessment of Real Estate Collateralized Financial Products." Journal of Computer Science and Frontier Technologies 3.1 (2026): 127-140.

Li, Zheyu, Yue Hao, and Peifan Zeng. "CLASNet: A Cognitive Load–Aware CNN-LSTM-Attention Framework for Supply Chain Demand Forecasting and Adaptive Human–Computer Interaction." Journal of Computer Science and Frontier Technologies 3.2 (2026): 44-57.

Yao, Yilin, Wenchao Zhang, and Mengtong Li. "Cloud-Edge Federated Incremental Learning Framework for PMSM Efficiency Optimization with Lightweight CNN-LSTM Models and OTA Differential Deployment." Journal of Computer Science and Frontier Technologies 3.1 (2026): 98-111.

Gao, Ziyang, Yuanliang Qu, and Yi Han. "Cross-Lingual Sponsored Search via Dual-Encoder and Graph Neural Networks for Context-Aware Query Translation in Advertising Platforms." arXiv preprint arXiv:2510.22957 (2025).

Zhang, Yangyang, et al. "Cognitive Digital Twin Framework: Modeling and Real-Time Decision Making."

© 2026 The Author(s). Licensed under CC BY 4.0

ICLR 2026 the 2nd Workshop on World Models: Understanding, Modelling and Scaling. 2026.

Xu, Chengde, Chongye Wang, and Zeyu Li. "H2-LBM: A Hierarchical Hybrid Deep Reinforcement Learning Framework for L7 Load Balancing and Global Traffic Scheduling in Multi-Cloud LLM Serving." Fundamental Scientific Reports in Multidisciplinary Areas 2.02 (2026): 109-120.

Zhang, Jingsi, and Jixiang Ding. "CA-PPO-AXI: Coverage-Aware Deep Reinforcement Learning for Adaptive AXI Bus Protocol Verification." Fundamental Scientific Reports in Multidisciplinary Areas 2.02 (2026): 159-169.

Hua, Zonghui, Yue Luo, and Yunting Ling. "LPG-SAC-TCO: A Large Language Model--Policy-Guided Soft Actor-Critic Framework for Total Cost of Ownership--Optimal Energy Management of Hydrogen Fuel Cell Heavy-Duty Trucks." European Journal of AI, Computing & Informatics 2.3 (2026): 1-11.

© 2026 The Author(s). Licensed under CC BY 4.0