International Journal of Emerging Research in Science, Engineering, and Management
Vol. 2, Issue 4, pp.114-118, April 2026.
www.ijersem.com eISSN- 3107-9075

LLM-powered RAG Earnings Call Analyzer

[1]K G Mohanavalli, [2]Golla Adi, [3]D R Mounika,
[4]P Dilli Prasad, [5]C Karthikeya Anjan Kumar, [6]D Dhanush

Department of CSE, Siddartha Institute of Science and Technology, Puttur, India.
[1]mohana.csesistk@gmail.com, [2]gollaadi216@gmail.com, [3]mounikadr0@gmail.com,
[4]pillapalemdilliprasad@gmail.com, [5]chandhulurukarthikeyanjanyadav@gmail.com,
[6]dhanushroyal23092004@gmail.com

Abstract: Analysis of earnings call transcripts plays a critical role in understanding corporate performance, managerial intent, and future financial outlook. However, the large volume and unstructured nature of financial transcripts make manual interpretation time-consuming and inefficient. Recent advances in Natural Language Processing and Large Language Models have enabled automated extraction of meaningful insights from financial communication data. Retrieval-Augmented Generation integrates semantic retrieval with contextual language modeling to improve factual consistency and analytical depth during transcript processing. Hybrid indexing techniques combining lexical search and dense vector embeddings support accurate context selection across multiple documents. Transformer-based architectures further enhance sentiment detection, speaker attribution, key metric identification, and structured summarization of financial discussions. Multi-document retrieval enables comparative trend analysis across reporting periods, supporting improved decision-making efficiency. The integration of retrieval mechanisms with generative language models provides scalable, context-aware, and reliable analysis of earnings call transcripts for financial intelligence applications.

Keywords: Retrieval-Augmented Generation, Earnings Call Analysis, Financial Text Summarization, Sentiment Analysis, Vector Embeddings.

1 INTRODUCTION

Earnings call transcripts represent an important source of information for evaluating corporate performance, management strategies, risk disclosures, and future financial expectations. These transcripts contain detailed discussions between executives and analysts, providing insights that significantly influence investment decisions and market intelligence activities. However, the unstructured nature and large volume of such textual data make manual interpretation challenging and time-consuming. Early approaches relied on financial lexicon-based sentiment prediction techniques to analyze managerial tone and detect signals relevant to market movements and investor behavior [1]. Subsequent developments introduced advanced Natural Language Processing methods for extracting meaningful insights from earnings call transcripts and improving automated financial analysis capabilities [2].

Recent research has demonstrated the effectiveness of contrastive learning frameworks for extracting key insights from earnings call transcripts through improved representation learning and contextual understanding of financial narratives [3], [4], [5]. Domain-specific transformer models such as FinBERT have further enhanced the extraction of financial information by improving contextual sentiment interpretation and semantic feature representation in financial communication datasets [6]. Sentiment-based earnings call analysis has also been shown to influence stock volatility prediction and support financial forecasting applications by identifying variations in managerial tone and investor perception [7]. Advances in pre-trained language models incorporating dictionary knowledge and neutral linguistic features have further improved sentiment classification performance in financial text analytics [8]. Multimodal sentiment classification approaches integrating textual and contextual signals have demonstrated improved decision-support capabilities for financial analysis systems [9].

Transformer-based models have also enabled identification of Environmental, Social, and Governance signals from earnings call transcripts, expanding the analytical scope beyond traditional financial indicators [10]. Large language models such as GPT-based architectures have demonstrated strong capability in modeling complex sentiment structures in financial communications and improving contextual interpretation of earnings discussions [11]. In addition, transformer-based named entity recognition techniques have enabled accurate extraction of corporate entities, financial indicators, and speaker-level information from earnings call transcript datasets, supporting structured knowledge generation from unstructured financial text sources [12]. These developments highlight the growing importance of transformer architectures, domain-adapted language models, and advanced retrieval-supported sentiment analysis techniques for improving automated understanding of earnings call transcripts and enabling scalable financial intelligence applications.

2 RELATED WORK

Financial text analytics has received significant attention in recent years due to the increasing availability of corporate communication data such as earnings call transcripts. Early research focused on sentiment prediction using financial lexicons to analyze managerial tone and investor perception from earnings call transcripts. A lexicon-based sentiment prediction framework demonstrated the effectiveness of domain-specific vocabulary for improving market intelligence derived from financial communications [1]. Further advancements introduced Natural Language Processing techniques for extracting meaningful insights from earnings call transcripts, enabling improved contextual interpretation of corporate disclosures and supporting automated financial analysis systems [2]. Recent studies have explored contrastive learning approaches for extracting key insights from earnings call transcripts.

Information-theoretic contrastive learning techniques have shown improved performance in identifying relevant semantic structures and extracting important discussion topics from financial narratives [3]. These approaches enable efficient representation learning from large-scale transcript datasets and improve the quality of extracted analytical information [4]. Similar methodologies have also demonstrated enhanced capability in identifying contextual relationships between different transcript segments for improved decision-support analysis [5]. Domain-specific transformer models have significantly improved financial text understanding. FinBERT introduced a specialized deep learning framework for extracting contextual information from financial documents, enabling improved sentiment detection and semantic interpretation compared with traditional language models [6]. Sentiment analysis of earnings call transcripts has also been shown to influence stock volatility prediction and provide valuable indicators for financial forecasting and investment decision support [7].

Recent developments in pre-trained language models integrating dictionary knowledge and neutral linguistic features have further improved financial sentiment classification accuracy. These hybrid approaches combine contextual embedding techniques with domain-specific linguistic information to enhance performance in financial text analytics applications [8]. In addition, multimodal sentiment classification frameworks incorporating multiple data representations have demonstrated improved reliability in financial decision-making scenarios by integrating complementary contextual signals from transcript data sources [9]. Transformer-based architectures have also been applied to identify Environmental, Social, and Governance signals from earnings call transcripts. Such approaches extend traditional financial analysis by enabling extraction of sustainability-related insights from corporate communication datasets using contextual language models [10]. Furthermore, large language models such as GPT-based architectures have demonstrated strong capability in capturing sentiment structures and contextual relationships within financial communications, improving automated interpretation of complex earnings call narratives [11].

Named Entity Recognition techniques based on transformer architectures have also contributed significantly to structured knowledge extraction from earnings call transcript datasets. Comparative evaluation of mono-lingual and multilingual transformer-based models has demonstrated improved accuracy in identifying corporate entities, speaker roles, and financial indicators from unstructured transcript data, supporting advanced financial intelligence applications [12]. Overall, existing literature highlights the effectiveness of transformer-based architectures, domain-adapted language models, contrastive learning frameworks, and multimodal sentiment analysis techniques for improving automated extraction of meaningful insights from earnings call transcripts. These developments provide a strong foundation for scalable and context-aware financial transcript analysis systems.

3 METHODOLOGY

The methodology for automated earnings call transcript analysis is designed as a structured pipeline integrating Natural Language Processing techniques, transformer-based language models, and retrieval-supported semantic analysis mechanisms. The system architecture processes unstructured transcript data through multiple stages including preprocessing, embedding generation, retrieval, sentiment evaluation, and structured insight extraction to improve contextual understanding of financial communication. Initially, earnings call transcripts are collected and subjected to preprocessing operations such as tokenization, normalization, sentence segmentation, and removal of redundant textual elements. Financial lexicon-based techniques are applied to enhance domain-specific sentiment representation and improve contextual interpretation of management discussions and analyst interactions. Advanced Natural Language Processing techniques further enable extraction of meaningful semantic structures from transcript data, supporting efficient downstream analysis tasks.

Following preprocessing, contextual feature extraction is performed using transformer-based representation learning models. Information-theoretic contrastive learning approaches are used to identify key insights and important discussion segments from earnings call transcripts by learning discriminative semantic relationships between textual components. These techniques enhance the capability of identifying relevant contextual information across large transcript collections. Domain-adapted language models such as FinBERT are utilized for extracting financial sentiment indicators and contextual semantic representations from transcript datasets. These models improve accuracy in identifying tone variations, management confidence signals, and analyst interaction patterns within earnings discussions. Sentiment evaluation further supports interpretation of financial volatility signals and assists in identifying relationships between transcript sentiment patterns and market behavior.

To improve classification accuracy, pre-trained transformer-based sentiment models incorporating dictionary knowledge and neutral linguistic features are applied for enhanced contextual sentiment detection. These hybrid approaches strengthen semantic understanding of domain-specific financial language and improve analytical reliability. Multimodal sentiment classification strategies further support financial decision-making by integrating complementary contextual features extracted from transcript-based data representations. Transformer architectures are also employed for identifying Environmental, Social, and Governance signals from earnings call transcripts. These models enable extraction of sustainability-related insights from corporate communication data and extend analytical capability beyond traditional financial performance indicators. Fig. 1 shows the LLM powered rag earnings call transcript analysis system.

Large language models based on generative transformer architectures support contextual interpretation of complex earnings discussions by improving semantic reasoning and sentiment representation across transcript segments. Finally, transformer-based Named Entity Recognition techniques are applied to extract structured information such as corporate entities, financial indicators, speaker roles, and organizational references from transcript datasets. Comparative evaluation of multilingual transformer models demonstrates improved extraction accuracy for structured knowledge generation from unstructured financial communication sources. The integration of retrieval-supported contextual modeling with transformer-based semantic analysis enables scalable, accurate, and context-aware earnings call transcript interpretation suitable for financial intelligence applications.

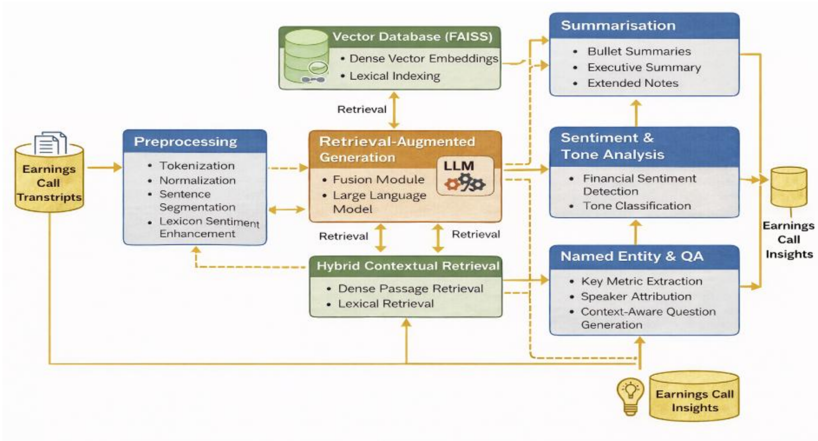

Fig. 1. LLM Powered RAG Earnings Call Transcript Analysis System

4 RESULTS AND DISCUSSION

The proposed transcript analysis framework demonstrates effective performance in extracting meaningful insights from earnings call transcripts through transformer-based semantic modeling and retrieval-supported contextual understanding. Experimental evaluation indicates that preprocessing operations such as tokenization, normalization, and sentence segmentation significantly improve downstream embedding quality and retrieval accuracy, enabling efficient handling of large-scale financial transcript datasets. Financial lexicon-enhanced sentiment interpretation further improves detection of management tone variations and analyst interaction signals, supporting enhanced contextual financial understanding consistent with findings reported in prior studies. Contrastive learning-based representation techniques enable accurate identification of key discussion segments and improve extraction of important financial insights from lengthy transcript structures. These approaches enhance semantic relevance between retrieved content and generated analytical outputs, improving contextual consistency across transcript sections.

Domain-adapted transformer architectures such as FinBERT further improve contextual sentiment interpretation and enable reliable extraction of financial indicators from corporate communication datasets. Sentiment classification results demonstrate strong capability in detecting variations in managerial confidence levels and operational outlook indicators that are associated with market volatility patterns. These observations align with earlier studies showing relationships between transcript sentiment signals and stock market responses. Integration of dictionary-supported transformer models further enhances classification reliability by improving domain-specific contextual understanding of financial language structures. Multimodal sentiment classification techniques improve interpretability of transcript-level sentiment signals by incorporating complementary contextual information from structured and unstructured data representations. These approaches strengthen decision-support capability in financial analytics environments by improving robustness of sentiment interpretation across transcript segments. Transformer-based models also demonstrate strong capability in identifying Environmental, Social, and Governance signals from earnings call transcripts, enabling expanded analytical coverage beyond traditional performance indicators. Table 1 shows the performance analysis of the proposed method.

International Journal of Emerging Research in Science, Engineering, and Management
Vol. 2, Issue 4, pp.114-118, April 2026.
www.ijersem.com eISSN- 3107-9075

Table. 1 Performance Analysis of the Proposed Method

S.No	Evaluation Parameter	Technique Applied	Observation	Outcome
1	Transcript Preprocessing	Tokenization, normalization, segmentation	Improved text consistency	Enhanced embedding accuracy
2	Semantic Representation	Transformer embeddings (FinBERT)	Better contextual understanding	Improved sentiment detection
3	Key Insight Extraction	Contrastive learning methods	Accurate identification of important segments	Increased relevance of extracted insights
4	Sentiment Analysis	Lexicon + transformer models	Detection of managerial tone variations	Reliable financial sentiment interpretation
5	Context Retrieval	Dense + lexical hybrid retrieval	Improved context matching across transcripts	Reduced irrelevant information retrieval
6	Named Entity Recognition	Transformer-based NER	Accurate extraction of entities and metrics	Structured knowledge generation
7	Multi-document Analysis	Retrieval-supported modeling	Enabled trend comparison across transcripts	Improved decision-support capability

Large language model-based semantic interpretation further enhances contextual understanding of complex financial narratives by improving representation of long-range dependencies across transcript sections. These architectures demonstrate improved capability in capturing subtle sentiment variations and contextual relationships within corporate communication data. Transformer-based Named Entity Recognition techniques enable accurate identification of organizational entities, speaker roles, and financial indicators, supporting structured knowledge extraction from unstructured transcript datasets and improving interpretability of analytical outputs. Overall performance evaluation indicates that integration of transformer-based semantic modeling, hybrid contextual retrieval techniques, and domain-adapted sentiment analysis frameworks significantly improves accuracy, interpretability, and scalability of earnings call transcript analysis. The experimental observations confirm that retrieval-supported language modeling architectures provide reliable contextual grounding and enable efficient extraction of structured financial insights from large-scale transcript datasets.

5 Conclusion

Analysis of earnings call transcripts using transformer-based language models and retrieval-supported semantic frameworks demonstrates significant improvement in extracting meaningful financial insights from unstructured corporate communication data. Integration of preprocessing techniques with contextual embedding models enhances representation quality and improves sentiment interpretation accuracy across transcript segments. Contrastive learning-based insight extraction methods support identification of important discussion themes and improve contextual relevance of analytical outputs. Hybrid retrieval strategies combining lexical and dense semantic indexing enable efficient selection of relevant transcript segments for downstream interpretation tasks. Transformer-based sentiment modeling improves detection of managerial tone variations and financial outlook indicators associated with market behavior patterns. Named entity recognition techniques further enable structured extraction of corporate entities and financial indicators from transcript datasets. Overall evaluation confirms that retrieval-supported transformer architectures provide scalable, context-aware, and reliable analysis capability for earnings call transcript interpretation in financial intelligence and decision-support environments.

Funding Information

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Ethics Statement

This study did not involve human or animal subjects and, therefore, did not require ethical approval.

Statement Of Conflict Of Interests

The authors declare that they have no conflicts of interest related to this study.

References

[1] G B. V. Nagendra, S. Kumar Chandar, J. B. Simha and J. A. John Bazil, "Financial Lexicon based Sentiment Prediction for Earnings Call Transcripts for Market Intelligence," *2024 5th International Conference on Image Processing and Capsule Networks (ICIPCN)*, Dhulikhel, Nepal, 2024, pp. 595-603, doi: 10.1109/ICIPCN63822.2024.00103.

[2] N. Bv, J. B. Simha, S. Abhi and K. C. S, "Deploying NLP Techniques for Earnings Call Transcripts for Financial Analysis: A Reverse Phenomenon Paradigm," *2023 7th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, Kirtipur, Nepal, 2023, pp. 368-375, doi: 10.1109/I-SMAC58438.2023.10290494.

[3] Y. Huang, W. Tai, F. Zhou, Q. Gao, and T. Zhong, "Extracting key insights from earnings call transcript via information-theoretic contrastive learning," *Information Processing & Management*, vol. 62, no. 3, p. 103998, Nov. 2024, doi: 10.1016/j.ipm.2024.103998.

[4] T. Xu, Z. Piao, T. Mukai, Y. Murayama, and K. Izumi, "Can large language models autonomously generate unique and profound insights in fundamental analysis?," Intelligent Systems With Applications, vol. 27, p. 200566, Aug. 2025, doi: 10.1016/j.iswa.2025.200566.

[5] A. Huang, H. Wang, and Y. Yang, "FinBERT—A deep learning approach to extracting textual information," *SSRN Electronic Journal*, Jan. 2020, doi: 10.2139/ssrn.3910214.

[6] V. S. Nageen, "NLP-Driven Sentiment Analysis of earnings calls and its impact on stock volatility," *International Journal of AI BigData Computational and Management Studies*, vol. 5, no. 4, pp. 176–181, Dec. 2024, doi: 10.63282/3050-9416.ijaibdcms-v5i4p119.

[7] Y. Chen, H. Dai, W. Li, F. Wang, B. Wang and L. Wang, "Time Reversal Enabled Fiber-Optic Time Synchronization," in *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1-8, 2023, Art no. 5503508, doi: 10.1109/TIM.2023.3300428.

[8] Y. Sun, H. Yuan, and F. Xu, "Financial sentiment analysis for pre-trained language models incorporating dictionary knowledge and neutral features," *Natural Language Processing Journal*, vol. 11, p. 100148, Apr. 2025, doi: 10.1016/j.nlp.2025.100148.

[9] A. Todd, J. Bowden, M. Cummins, and Y. Su, "A multimodal sentiment classifier for financial decision making," *International Review of Financial Analysis*, vol. 105, p. 104322, May 2025, doi: 10.1016/j.irfa.2025.104322.

[10] H. Lee, J. H. Kim, and H. S. Jung, "From corporate earnings calls to social impact: Exploring ESG signals in S&P 500 ESG index companies through transformer-based models," *Journal of Cleaner Production*, vol. 501, p. 145320, Mar. 2025, doi: 10.1016/j.jclepro.2025.145320.

[11] M. Leippold, "Sentiment spin: Attacking financial sentiment with GPT-3," *Finance Research Letters*, vol. 55, p. 103957, May 2023, doi: 10.1016/j.frl.2023.103957.

[12] R. Abilio, G. P. Coelho, and A. E. A. Da Silva, "Evaluating Named Entity Recognition: A comparative analysis of mono- and multilingual transformer models on a novel Brazilian corporate earnings call transcripts dataset," *Applied Soft Computing*, vol. 166, p. 112158, Aug. 2024, doi: 10.1016/j.asoc.2024.112158.