

**ENHANCING RETRIEVAL-AUGMENTED GENERATION FOR DATA
SCIENCE: A COMPREHENSIVE FRAMEWORK FOR ACADEMIC
LITERATURE NAVIGATION**

by
AHMET YASIN AY TAR

Submitted to the Graduate School of Engineering and Natural Sciences
in partial fulfilment of
the requirements for the degree of Master of Science

Sabancı University
December 2024

AHMET YASİN AY TAR 2024 ©

All Rights Reserved

ABSTRACT

ENHANCING RETRIEVAL-AUGMENTED GENERATION FOR DATA SCIENCE: A
COMPREHENSIVE FRAMEWORK FOR ACADEMIC LITERATURE NAVIGATION

AHMET YASIN AY TAR

Data Science, MSc. Thesis, December 2024

Thesis Supervisor: Assoc. Prof. Kemal Kılıç

Thesis Co-Supervisor: Assoc. Prof. Kamer Kaya

Keywords: Retrieval-Augmented Generation (RAG), Data Science, Literature Retrieval,
Academic Insights, Large Language Models (LLM)

In the rapidly evolving field of data science, efficiently navigating the expansive body of academic literature is crucial for informed decision-making and innovation. This thesis presents an enhanced Retrieval-Augmented Generation (RAG) application designed to assist data scientists in accessing precise and contextually relevant academic resources. The application integrates advanced techniques, including GeneRation Of Bibliographic Data (GROBID), fine-tuning embedding model, semantic chunking, and an abstract-first retrieval method, to significantly improve the relevance and accuracy of the retrieved information. A comprehensive evaluation using the Retrieval-Augmented Generation Assessment System (RAGAS) framework demonstrates substantial improvements in key metrics, particularly Context Relevance, underscoring the system's effectiveness in reducing information overload and enhancing decision-making processes. Our findings highlight the potential of this enhanced RAG system to transform academic exploration within data science, providing a valuable tool for researchers and practitioners alike.

ÖZET

VERİ BİLİMİ ODAKLI BİR RAG MODELİ: AKADEMİK LİTERATÜR TARAMA İÇİN KAPSAMLI BİR ARAÇ

AHMET YASIN AY TAR

Veri Bilimi, Yüksek Lisans Tezi, Aralık 2024

Tez Danışmanı: Assoc. Prof. Kemal Kılıç

Tez Ortak Danışmanı: Assoc. Prof. Kamer Kaya

Anahtar Kelimeler: Retrieval-Augmented Generation (RAG), Veri Bilimi, Literatür
Tarama, Akademik İlgörüler, Büyük Dil Modelleri (LLM)

Hızla gelişen veri bilimi alanında, oldukça kapsamlı olan akademik literatürü etkili bir şekilde araştırmak, karar alımı ve inovasyon için oldukça önemlidir. Bu makale, veri bilimcilerin doğru ve bağlamsal olarak ilgili akademik kaynaklara erişmelerine yardımcı olmak amacıyla tasarlanmış bir Retrieval-Augmented Generation (RAG) uygulamasını sunmaktadır. Uygulama, GeneRation Of Bibliographic Data (GROBID), embedding modelini fine-tune etme, semantic chunking ve hiyerarşik bilgi çekme yöntemi gibi ileri teknikleri entegre ederek getirilen bilgilerin alaka düzeyini önemli ölçüde artırmaktadır. Retrieval-Augmented Generation Assessment System (RAGAS) çerçevesi kullanılarak yapılan kapsamlı bir değerlendirme, özellikle Context Relevancy (Bağlam Uyumu) metriği açısından önemli iyileştirmeler göstererek sistemin bilgi yükünü azaltma ve karar alma süreçlerini geliştirme konusundaki etkinliğini vurgulamaktadır. Bulgularımız, bu geliştirilmiş RAG sisteminin veri bilimi alanında akademik literatür taramasını dönüştürme potansiyelini ve araştırmacılar ile uygulayıcılar için değerli bir araç olma özelliğini ortaya koymaktadır.

ACKNOWLEDGEMENTS

I would like to extend my heartfelt gratitude to my thesis advisors, Dr. Kemal Kılıç and Dr. Kamer Kaya, for their unwavering guidance, support, and encouragement throughout the course of this research. Their insights and expertise have been invaluable in shaping this work.

I am also deeply thankful to Dr. Pınar Karagöz, Dr. Murat Kaya, and Dr. Yücel Saygın for serving as jury members.

Lastly, I would like to express my profound appreciation to my family for their constant support, patience, and encouragement, which have been my foundation throughout this journey.

Dedicated to my beloved family

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
1. INTRODUCTION	1
2. LITERATURE REVIEW	5
2.1. The Role and Challenges of Academic Guidance in Data Science	5
2.2. Retrieval-Augmented Generation as a Solution for Data Science	6
2.3. Applications of RAG in Various Domains	6
2.4. Challenges and Enhancements of RAG Systems	7
3. METHODS	11
3.1. Overview of RAG Architecture and Data Preparation	11
3.2. 5-Staged Enhancements	14
3.2.1. Data Cleaning using GROBID	14
3.2.2. Fine-tuning	15
3.2.3. Semantic Chunking	17
3.2.4. Abstract-first Method	18
3.2.5. Enhanced Prompting Technique	19
3.3. Evaluation Using RAGAS	20
4. RESULTS	23
4.1. Performance Metrics for Different Fine-tuning Experiments	23
4.2. Performance Metrics for RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 3	25
4.3. Performance Metrics for RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 6	28
5. DISCUSSION & CONCLUSIONS	30

BIBLIOGRAPHY.....	33
APPENDIX	37

LIST OF TABLES

Table 4.1. Performance Metrics for Different Fine-tuning Experiments. (FT = Fine-tuning, TB = Textbooks, G = GROBID, SNS = Semantic Node Splitter).....	23
Table 4.2. Significant Pairwise Comparisons for Context Relevance Using Tukey’s HSD Test.....	24
Table 4.3. Performance Metrics for Different RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 3. (B = Baseline, G = GROBID, FT = Fine-tuning, SC = Semantic Chunking, AF = Abstract First, EPT = Enhanced Prompting Technique).....	26
Table 4.4. Significant Pairwise Comparisons for Different Metrics Using Tukey’s HSD Test for Table 4.3. (CR = Context Relevance, F = Faithfulness, AR = Answer Relevance)	27
Table 4.5. Performance Metrics for Different RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 6. (B = Baseline, G = GROBID, FT = Fine-tuning, SC = Semantic Chunking, AF = Abstract First, EPT = Enhanced Prompting Technique).....	28
Table 4.6. Significant Pairwise Comparisons for Different Metrics Using Tukey’s HSD Test for Table 4.5. (CR = Context Relevance, F = Faithfulness, AR = Answer Relevance)	29

LIST OF FIGURES

Figure 3.1. Original RAG Architecture.....	12
Figure 3.2. Our Enhanced RAG Architecture	12
Figure 3.3. Three Performance Metrics of RAGAS Framework and Their Cor- responding Aspects	22

1. INTRODUCTION

Like other Generative Artificial Intelligence (GenAI) technologies—such as image, video, and audio generation—Large Language Models (LLMs) that produce text also possess the potential to disrupt both business and social life by transforming how information is processed and decisions are made. However, a significant drawback of LLMs is their current tendency to produce "hallucinations" (i.e., inaccurate or misleading outputs), which can undermine trust in their applications. To address this issue, ongoing research explores Retrieval-Augmented Generation (RAG) as a promising solution to mitigate these inaccuracies. Moreover, across various disciplines, practitioners, and researchers are actively seeking to harness the capabilities of LLMs, particularly as intelligent assistants to enhance decision-making processes. RAG can potentially improve the effectiveness of intelligent assistants by integrating relevant information from external sources, allowing for more accurate and contextually appropriate responses. Data science, as a critical field in digital transformation that relies on data-driven insights, is no exception; it stands to benefit from advancements in LLM technology while addressing its inherent weaknesses.

Data science is an exploratory process that necessitates the formulation and validation of hypotheses using historical datasets. The models developed by data scientists essentially represent these hypotheses, and creating effective models relies on the formulation of pertinent questions. Therefore, data scientists benefit from exposure to various alternatives, and academic literature serves as a vital resource that enhances their creative thinking abilities. However, as data science continuously evolves, it generates an overwhelming volume of research, making the efficient access to and utilization of academic literature increasingly challenging. This surge in available information complicates the ability of researchers and practitioners to identify and integrate relevant insights into their work, highlighting the need for effective tools and strategies to navigate this expanding landscape.

To mitigate this issue, in this research, we have developed a RAG application specifically designed to aid data scientists in navigating academic resources and obtaining guidance

tailored to their unique needs. This application encompasses a broad spectrum of data science subfields, including machine learning, deep learning, time series modeling, synthetic data generation, reinforcement learning, data preprocessing, feature engineering, and evaluation. The value of our RAG application is evident in its ability to provide precise and contextually relevant responses to specific queries posed by data scientists. This tool not only conserves time but also improves the accuracy of decisions made throughout the data science workflow by integrating advanced techniques for retrieving and embedding pertinent academic insights into a LLM. Consequently, the application enables users to make well-informed decisions regarding methodologies and techniques, thereby contributing to the development of more effective and innovative solutions for their projects.

Our objective with this application is to systematically combine and assess the performance of a range of established techniques within the RAG framework, reporting their effectiveness within a particular domain to generate insights valuable to researchers. By conducting a comprehensive evaluation of these methods in a real-world setting, we aim to facilitate the broader adoption of these approaches in both academic and industrial contexts.

Our RAG application responds to a significant need within the data science community for accessible and relevant academic guidance. Unlike traditional literature review methods, which can be time-consuming and overwhelming due to the vast amount of available information, our application facilitates the process by extracting and embedding only the most pertinent sections of articles into the LLM. This approach accelerates decision-making and ensures that the information provided is both current and directly relevant to the user's query, thereby addressing the prevalent issue of information overload in data science research.

The primary innovation of our RAG application lies in its *five-stage enhancement process*, which substantially improves the traditional RAG pipeline. Initially, we utilize GROBID (GeneRation Of Bibliographic Data) to clean and structure article data before constructing the vector database, ensuring accuracy and organization of the information (Lopez, 2024). Next, we implement a specialized fine-tuning process using academic textbooks, which enhances the embedding model's capacity to interpret and process complex queries. Furthermore, our application employs semantic chunking rather than traditional recursive chunking which ensures that text is divided into meaningful units and results in a more coherent and relevant vector database (Kamradt, 2024). We also adopt an abstract-first method, prioritizing searches within a separate vector database of article abstracts, thereby streamlining the retrieval process. Finally, advanced prompting techniques are

employed to optimize the LLM’s responses, leading to more accurate and contextually appropriate answers. Llamaindex and LangChain were instrumental in creating fine-tuning data and RAG pipelines, respectively (Liu, 2022; Chase, 2022). Additionally, OpenAI’s GPT-4o model was employed as the LLM throughout the study, chosen for its cost-effectiveness and efficiency.

To validate the effectiveness of our application, we employed the RAGAS (Retrieval Augmented Generation for Academic Search) framework and conducted tests with 50 sample questions covering various domains and problem types (Es et al., 2023). This rigorous testing process confirmed the application’s capability to deliver high-quality, relevant responses, highlighting its potential as a useful tool for data scientists. Although all metrics in the RAGAS framework—Context Relevance, Faithfulness, and Answer Relevance—are important, our primary focus was on improving Context Relevance. This metric serves as the main indicator of the system’s ability to retrieve pertinent academic content, which aligns with the central objective of this study: enhancing the retrieval process for relevant information in academic literature for data scientists. The other metrics, such as Faithfulness and Answer Relevance, assess the quality of answers generated using the retrieved context, but since the language model used for answer generation remained unchanged, these metrics are less indicative of improvements brought by the fine-tuning process and other retrieval enhancements. Therefore, Context Relevance serves as the most appropriate measure to demonstrate advancements in retrieval capabilities across different configurations.

Research Objectives:

- To develop an enhanced Retrieval-Augmented Generation (RAG) system tailored to the needs of data scientists, incorporating advanced techniques such as data cleaning, fine-tuned embedding models, semantic chunking, different retrieval mechanisms, prompt engineering, and systematic evaluation of each component’s impact on performance using key metrics within the RAGAS framework.
- To conduct a focused analysis of the fine-tuning process, characterized by its experimental flexibility and various potential approaches, involving the exploration of alternative configurations through adjustments in data sources and preprocessing techniques, to assess how different fine-tuning strategies influence key metrics.
- To evaluate the performance of the enhanced RAG system using the RAGAS framework, focusing on metrics such as context relevancy, faithfulness, and answer relevance.

- To design and utilize 50 questions across different subdomains of data science to rigorously test the system’s effectiveness and applicability.
- To demonstrate significant improvements in key metrics when compared to baseline performance, showing the effectiveness of the proposed enhancements.

The remainder of this thesis is organized as follows. Chapter 2 provides the Literature Review, presenting an overview of existing research on academic literature retrieval systems, with a focus on advancements in Retrieval-Augmented Generation (RAG) models. Chapter 3, the Methods section, details the five-stage enhancement process of the RAG pipeline, including data cleaning, fine-tuning of embedding models, semantic chunking, abstract-first retrieval, and prompting strategies, as well as the evaluation using the RAGAS framework. Chapter 4 presents the Results, discussing the performance of the enhanced RAG system and comparing it to baseline models. Finally, Chapter 5 combines Discussion & Conclusions, summarizing key contributions, interpreting the results, and suggesting potential areas for future research.

2. LITERATURE REVIEW

2.1 The Role and Challenges of Academic Guidance in Data Science

Academic guidance is an essential element in the work of a data scientist, providing a framework for grasping existing research and directing new, creative project directions. Engaging with scholarly articles allows data scientists to identify established methodologies, discern gaps in current research, and draw connections between their specific challenges and those addressed in prior studies. This process is essential for both refining research questions and selecting or adjusting methods for data analysis and model development. As noted by Mysore et al. (2023), data scientists heavily depend on the literature to assess what has been accomplished, determine whether similar problems have been addressed, and find solutions to related challenges.

In recent years, the volume of literature on AI and machine learning has expanded rapidly, with publications doubling approximately every 23 months as of 2022 (Krenn et al., 2022). This continual growth has resulted in an overwhelming amount of information for data scientists to sift through. The difficulty of pinpointing the most relevant and high-quality articles amid the vast number of publications adds complexity to the task. Consequently, data scientists often rely on keyword searches or similar techniques, which may be inadequate due to the difficulty in crafting effective search queries. They may struggle to identify the appropriate terms or strategies to locate relevant articles efficiently and promptly, leading to potential gaps in their literature reviews (Mysore et al., 2023).

2.2 Retrieval-Augmented Generation as a Solution for Data Science

A promising approach to addressing the challenge of navigating the extensive literature in data science is Retrieval-Augmented Generation (RAG) (Lewis et al., 2020). RAG’s capacity to access external knowledge bases makes it a valuable tool for managing and comprehending the vast literature in the field. The integration of Large Language Models (LLMs) within the RAG pipeline further enhances this capability by making the process of understanding complex academic papers and their core contributions more manageable. By retrieving contextually relevant information and summarizing it effectively, RAG assists data scientists in quickly grasping the key contributions of numerous papers, thus mitigating the burden of information overload. RAG models have demonstrated considerable promise in tasks such as question-answering and dialogue systems, where the ability to dynamically retrieve pertinent information from external sources is essential (Izacard and Grave, 2020). This capability is particularly valuable in the fast-growing field of data science, where synthesizing and analyzing an expanding body of literature is critical.

2.3 Applications of RAG in Various Domains

RAG has been adopted across various industries, including healthcare, legal, and education to enhance domain-specific applications and address the challenges of these fields. In healthcare, RAG models support clinicians by providing evidence-based recommendations from up-to-date medical literature and databases, enhancing diagnostic accuracy and treatment planning. For example, the MKRAG framework has improved medical question answering by integrating medical facts from external knowledge bases into the generation process, thereby increasing the accuracy of responses and aiding clinicians in making informed decisions (Shi et al., 2023b). In the legal field, one method to improve question-answering systems involves the integration of RAG models with Case-Based Reasoning (CBR). This CBR-RAG approach enhances LLM queries with contextually relevant legal cases, improving the retrieval of pertinent precedents and statutes. This method is among the strategies developed to provide more accurate and contextually relevant answers to complex legal queries, thereby increasing the efficiency and accuracy

of legal research (Wiratunga et al., 2024). Additionally, systems like LegalAsst utilize AI-powered tools for precise legal case retrieval, incorporating structured representations and decision-tree-based judgments to streamline court procedures, thus enhancing legal assistance and court productivity through improved retrieval accuracy of similar cases and relevant legislation (Han et al., 2024). In education, RAG models are utilized to improve math question-answering by incorporating vetted external content, such as scripts from high-quality open-source textbooks, into LLM prompts. This method has shown to improve response quality, assisting students and educators in navigating complex problems with contextually relevant support, although designers must balance the trade-offs between aligning responses with educational resources and meeting student preferences (Levonian et al., 2023).

RAG applications have significantly streamlined the academic literature review process. For example, KNIMEZoBot (Alshammari et al., 2023) integrates Zotero with KNIME and OpenAI’s language models to automate citation management and literature generation, reducing the time and effort required for comprehensive reviews. LitLLM (Agarwal et al., 2024) further enhances this process by synthesizing large volumes of research into concise, field-specific summaries, making it an invaluable tool for scientific researchers. In the biomedical field, RefAI (Li et al., 2024) employs a GPT-powered RAG system to manage complex terminologies and research details with precision. Collectively, these tools illustrate the growing influence of RAG technology in transforming literature reviews, offering tailored and efficient solutions across diverse academic disciplines.

2.4 Challenges and Enhancements of RAG Systems

Despite the advancements offered by RAG applications, substantial challenges remain, particularly regarding information overload and contextual relevance. For instance, current models often retrieve passages indiscriminately, without regard to whether they are factually grounded or contextually appropriate, which can result in low-quality outputs containing irrelevant or off-topic information (Shi et al., 2023a). Furthermore, the accuracy and reliability of the generated content are not guaranteed, as these models are not always explicitly trained to leverage facts from the retrieved passages, leading to potential inconsistencies (Gao et al., 2023). Additionally, handling complex and domain-specific queries remains a significant issue, as RAG models often require fine-tuning to improve

their accuracy and effectiveness in specialized fields (Siriwardhana et al., 2023). Moreover, obtaining a high-quality retrieval model that effectively fetches relevant documents is an ongoing challenge that impacts the overall performance of these systems (Arora et al., 2023).

To address the limitations identified in traditional Retrieval-Augmented Generation (RAG) models, we propose a comprehensive five-stage enhancement process that significantly improves the performance, accuracy, and relevance of RAG applications. This process is designed to handle the complexities of domain-specific queries, ensuring that the information retrieved is not only pertinent but also aligned with the user’s context and intent. Each stage of our approach contributes to refining different aspects of the RAG system, from data cleaning and embedding fine-tuning to semantic chunking, abstract-first retrieval, and advanced prompting techniques. Collectively, these innovations have the potential to significantly reduce information overload and improve the contextual relevance of responses, thereby providing a robust solution to the common challenges encountered by traditional RAG models.

A key component of our approach is the use of GROBID (Lopez, 2024), an open-source tool renowned for its accuracy in reference parsing and metadata extraction. GROBID is utilized in our application to structure and clean data from PDFs, ensuring the information is well-organized before being embedded into our vector database. A study evaluating multiple reference parsing tools demonstrated that GROBID outperformed its peers, achieving the highest F1 score of 0.89 among the ten tools tested (Tkaczyk et al., 2018). This high level of precision is crucial for maintaining the quality of metadata and ensuring that a RAG system delivers accurate and contextually appropriate responses. Moreover, GROBID’s success in various contexts, such as automated metadata extraction in the CERN Document Server and citation generation in CiteAssist, further underscores its versatility and reliability (Kaesberg et al., 2024; Boyd, 2015).

In the second stage of our approach, we emphasize the fine-tuning of the embedding model to enhance its capability to accurately retrieve domain-specific information, a critical step in improving the overall effectiveness of our RAG application. This fine-tuning process draws on methodologies from recent studies, which highlight the importance of adapting models to the specific contexts in which they are applied. For instance, Pu et al. (2024) demonstrated how domain-specific fine-tuning significantly improved retrieval accuracy for technical documentation in Electronic Design Automation (EDA) tools by benchmarking the performance of various models. Similarly, Nguyen et al. (2024) combined fine-tuning with iterative reasoning in a Q&A system, allowing for progressively refined answers, especially in complex scenarios. Natural Language Inference (NLI) to

fine-tune embedding models is employed in Dušek et al. (2023), enhancing their ability to understand nuanced domain-specific language, thereby improving the relevance and accuracy of retrieved content. Therefore, by incorporating these embedding fine-tuning techniques, we hypothesized that the RAG application would be better equipped to handle specialized and intricate queries common in data science workflows, ensuring that the information provided is both precise and contextually appropriate.

In the third stage of our approach, we apply semantic chunking to improve the coherence and relevance of the data used in our RAG application Kamradt (2024). This technique enables breaking down complex information into semantically meaningful units, which are then more effectively mapped to relevant knowledge sources. The application of semantic chunking has demonstrated significant usefulness in specific fields. For example, Wu et al. (2024) applied semantic chunking to organize medical knowledge into graph-based structures, ensuring the safe and contextually appropriate retrieval of medical information—a critical feature in healthcare where accuracy is essential. Additionally, Wu and Mooney (2022) utilized semantic chunking in Visual Question Answering (VQA), where it was used to decompose complex visual questions into simpler, semantically coherent parts. This approach significantly enhanced the model’s ability to interpret and respond accurately by linking these chunks to external knowledge sources. Consequently in the enhanced RAG architecture we considered integrating a similar semantic chunking technique to maintain high levels of contextual accuracy and relevance, ensuring that the information retrieved and presented is both coherent and directly applicable to the user’s query.

In the fourth stage of our enhancement process, we introduce an "abstract-first" method for RAG applications, which prioritizes searches within a separate vector database of article abstracts. This approach is intended to enhance the retrieval process by concentrating on the most concise and relevant sections of academic papers, which are the abstracts. The rationale behind this method is that abstracts typically provide a clear and succinct summary of the research, capturing the core findings and methodologies. Consequently, many existing datasets used for Named Entity Recognition (NER) models, which are essential for NLP tasks such as information extraction, concentrate solely on abstracts and preselected texts (Otto et al., 2023). This strategy mirrors the natural tendency of researchers to first read abstracts when assessing the relevance of a paper (Sambar et al., 2024). By targeting these sections first, a RAG application can quickly and efficiently determine the relevance of a document, thereby improving retrieval efficiency and reducing computational load. This approach is particularly advantageous in fields where quick access to respective information is critical, addressing the common challenge of exploring

large volumes of full-text content to identify key insights. The abstract-first method thus potentially represents a significant advancement in the design of RAG systems, offering a sophisticated alternative to traditional full-text search strategies.

In the fifth stage of our enhancement process, we implement advanced prompting techniques to optimize the responses generated by our RAG application. This stage involves incorporating strategies such as tip-offering and emotional prompting, which have been shown to enhance the performance and contextual relevance of Large Language Models (LLMs). For example, Salinas and Morstatter (2024) demonstrated that even minor adjustments to prompts, such as offering tips or slight rephrasing, can lead to substantial improvements in the precision and relevance of LLM outputs. This underscores the importance of carefully crafted prompts in refining the responses generated by RAG systems. Similarly, Li et al. (2023) showed that emotional prompting, where prompts are infused with emotional cues, can improve the LLM’s ability to produce more empathetic and contextually rich responses, which is particularly valuable in applications requiring a nuanced understanding of user needs, such as mental health support or customer service. By integrating these advanced prompting techniques into the RAG model, we aim to improve the performance of the LLM while also offering prompts that direct the model to produce more precise and contextually relevant responses, ultimately enhancing the utility and effectiveness of our application.

3. METHODS

In this section, we will first present the original RAG approach and provide the details of the enhanced RAG architecture we have developed in this research. Additionally, we will explain the experimental design and present the performance metrics used to evaluate the effectiveness of the alternatives.

3.1 Overview of RAG Architecture and Data Preparation

In the progression from an original RAG architecture to our enhanced version, significant modifications were made to improve the system’s accuracy, relevance, and efficiency. Figure 3.1 illustrates the basic RAG architecture, where a user’s query is first embedded and then used to search a vector database for relevant document chunks. These retrieved chunks, along with the original query, are then fed into a language model to generate the final output. In contrast, Figure 3.2 showcases our enhanced RAG architecture, where we have introduced several critical innovations, highlighted in blue, to address the limitations of the classical approach. These enhancements include advanced data cleaning, fine-tuned embedding models, semantic chunking, a novel abstract-first search strategy, and improved prompting techniques. The details of these improvements are outlined below.

To thoroughly evaluate our enhanced RAG application, we developed a comprehensive set of 50 questions covering various sectors such as telecommunications, agriculture, and others, reflecting the diverse challenges encountered in data science. These questions were meticulously crafted according to the CRISP-DM (Cross-Industry Standard Process for Data Mining) methodology (Chapman et al., 2000), a widely recognized framework in the data science community. The question set spans the three major stages of the CRISP-DM

framework: Data Preparation, Modeling, and Evaluation, with 10 questions dedicated to Data Preparation, 35 to Modeling, and 5 to Evaluation.

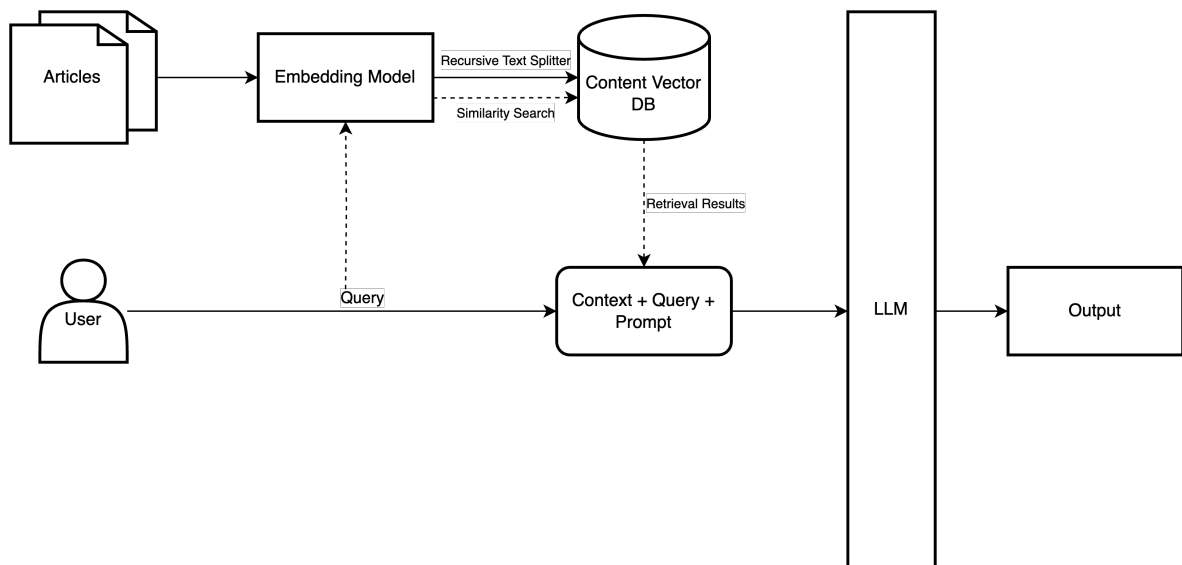


Figure 3.1 Original RAG Architecture

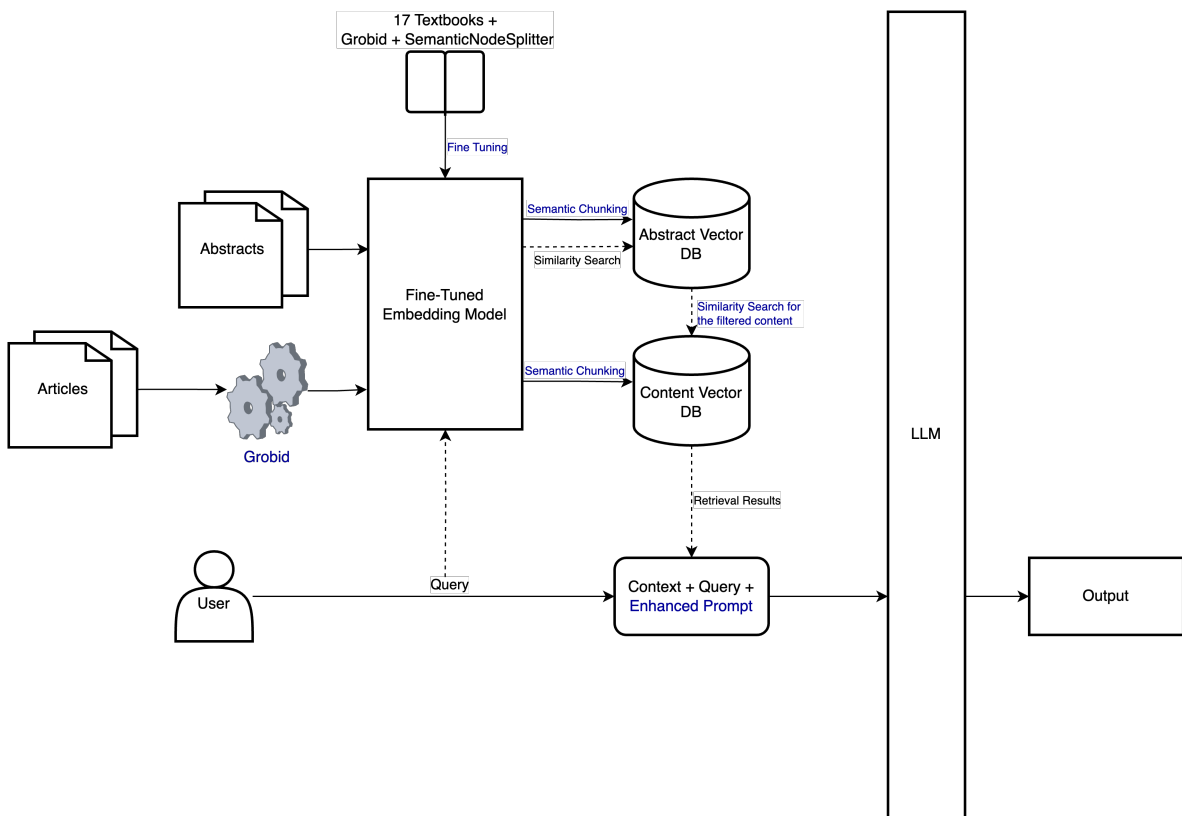


Figure 3.2 Our Enhanced RAG Architecture

Within the Data Preparation stage, we targeted essential areas such as feature selection (3 questions), outlier elimination (2 questions), and dimension reduction (5 questions). The Modeling stage, which often demands the most academic guidance, included questions on synthetic data generation (3 questions), classification (5 questions), regression (5 questions), clustering (5 questions), reinforcement learning (2 questions), image recognition (5 questions), natural language processing (5 questions), and time series analysis (5 questions). Lastly, the Evaluation stage comprised 5 questions focused on assessing model performance. Each question was designed to address specific subdomains and techniques, ensuring a rigorous assessment of the RAG application across multiple aspects of data science. The complete list of questions is provided in Appendix A.

For instance, one of our questions is: "As a data scientist in the financial sector tasked with predicting stock prices, I have used traditional feature selection methods like backward elimination. Which advanced feature selection technique should I consider to enhance model performance?" This format ensured that the questions were realistic and reflective of the actual challenges faced by professionals in the field.

In our analysis, arXiv API was employed to retrieve relevant articles necessary for constructing a vector database. We used arXiv (Wong, 2017) since it is a prominent platform that allows researchers to share their scientific and academic papers prior to undergoing the peer-review process (Khouya et al., 2024). For the analysis, each test question was first translated into a detailed search query, compatible with the arXiv API. This conversion process was facilitated by GPT-4o using a specific prompt designed to condense each question into a targeted search query: "Transform the following detailed inquiry into a concise and focused research query suitable for academic search databases. The summary should highlight the main topic, applicable domain, and specific requirements or goals in a maximum of 10 words." This approach ensured that each search term accurately represented the specific research objective associated with each question. Approximately 100 relevant articles were selected for each question, resulting in a dataset comprising roughly 5,000 articles. The content of the selected articles was downloaded as PDF files and subsequently used in the data-cleaning step. In parallel, the metadata for each article—such as publication date, title, authors, and abstracts—retrieved from the arXiv API, was also utilized. Following data cleaning, both the cleaned article content and the associated metadata were systematically integrated into the vector database. This process ensured that the database facilitated structured retrieval in accordance with academic standards.

3.2 5-Staged Enhancements

3.2.1 Data Cleaning using GROBID

To ensure a structured and accurate representation of the scientific documents sourced from arXiv, we employed GROBID (GeneRation Of Bibliographic Data), a widely-used machine learning tool that specializes in the extraction and parsing of PDFs (Lopez, 2024). GROBID converts unstructured PDF files into structured XML/TEI formats, facilitating subsequent text mining and information retrieval processes. The extraction of structured data from PDF documents is particularly important in the context of scientific publications, where figures, tables, and mathematical symbols can introduce noise into the data if not properly handled. By converting PDFs into a predictable and standardized format, GROBID ensures that the processed data retains the essential textual elements while discarding non-informative content, such as image annotations, complex mathematical notations, and extraneous spaces.

At its core, GROBID utilizes a series of machine learning-based sequence labeling models, arranged in a modular cascade, to parse different sections of a document. This process begins with the identification of document areas, including headers, footers, and body text, using layout-based features. Once the main sections are identified, GROBID applies more granular models to parse specific elements like author names, abstracts, and bibliographical references. This cascading approach allows the system to produce a highly structured output, including 55 distinct labels, making it particularly well-suited for scientific texts. Moreover, GROBID’s reliance on layout tokens—rather than just raw text—enables it to capture visual features such as font size, style attributes (e.g., bold or italic), and the relative positioning of text within the document. This information is crucial for correctly identifying structural elements like section titles, figure captions, and references that are often defined by their visual presentation.

GROBID’s training approach emphasizes the quality of its training data rather than relying on vast datasets, as seen in other systems like CERMINE (Tkaczyk et al., 2015). Instead of leveraging large sets derived from existing publisher XML documents, GROBID uses small, manually-labeled datasets generated directly from PDFs. This approach allows for more accurate alignment between noisy PDF content and the resulting structured

text, avoiding the inconsistencies often found in publisher-provided XML. Additionally, the system iterates on its training by continuously introducing error cases, enabling it to improve incrementally. Evaluation of GROBID is conducted using stable holdout sets from various publishers, ensuring realistic performance assessments across diverse document structures. However, the lack of diversity in some evaluation sets, mainly sourced from life sciences and preprints, remains a challenge, which the system addresses by progressively adding new datasets from different domains. Despite its smaller training size, GROBID consistently delivers accuracy comparable to larger models, making it highly efficient for scientific document parsing. A study evaluating multiple reference parsing tools demonstrated that GROBID outperformed its peers, achieving the highest F1 score of 0.89 among the ten tools tested (Tkaczyk et al., 2018)

By applying GROBID’s pipeline to the documents retrieved from arXiv, we were able to streamline the data ingestion process, transforming raw, unstructured PDF files into structured text suitable for embedding in our vector databases. This preprocessing step is foundational to the accuracy of subsequent stages in our enhanced retrieval-augmented generation pipeline, where the quality of document embeddings depends directly on the integrity and clarity of the input data.

3.2.2 Fine-tuning

The second stage focused on fine-tuning the embedding model to improve its ability to capture domain-specific nuances in order to ensure that both the articles and queries were embedded in a manner that accurately reflected the complexities of data science and academic literature. Given that the embedding model is essential for converting both queries and articles into vectors for storage in the vector database, it was imperative that the model be proficient in understanding complex academic language. We conducted a series of experiments to fine-tune the "sentence-transformers/all-mpnet-base-v2" model using a range of textbooks and various techniques Transformers (2021). These experiments included:

1. **Experiment 1:** Fine-tuning with 5 textbooks.
2. **Experiment 2:** Fine-tuning with 17 textbooks.
3. **Experiment 3:** Fine-tuning with 5 textbooks, incorporating GROBID for additional data cleaning.

4. **Experiment 4:** Fine-tuning with 17 textbooks, incorporating GROBID for additional data cleaning.
5. **Experiment 5:** Combining 17 textbooks, GROBID, and Semantic NodeSplitter.

The textbooks, listed in Appendix B and Appendix C, cover the same topics as our question set, ensuring that the fine-tuned models were well-aligned with the content they were intended to process. The fine-tuning process involved converting each textbook into nodes—coherent segments of text—and then generating question-answer pairs for each node using the OpenAI API. GROBID, specialized in processing academic content, was also employed to clean the fine-tuning dataset before creating these Q-A pairs. GROBID’s application to the textbooks resulted in well-prepared text for subsequent stages. As emphasized by Gunasekar et al. (2023), the importance of high-quality data in the context of fine-tuning cannot be overstated. This insight drove our decision to use GROBID for data cleaning, ensuring the integrity and relevance of the academic content used in our experiments (Lopez, 2024). Additionally, by incorporating the Semantic Node Splitter, which operates similarly to semantic chunking, we were able to segment nodes based on the semantic similarity of sentences rather than arbitrary text lengths (Kamradt, 2024). This approach allowed us to create nodes that were more semantically coherent instead of using preset token size to create nodes, which resulted in more meaningful Q-A pairs and a fine-tuned model better aligned with the complexities of academic content.

The number of Q-A pairs for each experiment is as follows: Experiment 1 produced 4,819 pairs, Experiment 2 resulted in 15,326 pairs, Experiment 3 generated 1,600 pairs, Experiment 4 created 4,460 pairs, and Experiment 5 yielded 4,132 pairs (altogether 30,337 pairs). As the number of textbooks increased between Experiments 1 and 2, the number of Q-A pairs grew substantially due to the greater volume of available data. However, when GROBID was applied in Experiments 3 and 4, a significant reduction in Q-A pairs was observed, as GROBID systematically removes extraneous content and focuses on the core textual material relevant for model fine-tuning. When the Semantic Node Splitter was introduced in Experiment 5, the number of Q-A pairs remained relatively consistent by adjusting the similarity threshold during node splitting to provide the optimal balance between quantity and the meaningfulness of the Q-A pairs. Since there is no established standard for the optimal number of Q-A pairs required for fine-tuning an embedding model, these varied experiments are designed to generate datasets of different sizes and structures for fine-tuning. This variability allows us to explore how different data formats influence the model’s performance and gain insights into the relationship between dataset size, content quality, and the resulting fine-tuned embeddings (Lopez, 2024).

3.2.3 Semantic Chunking

In constructing the vector database for our system, we employed semantic chunking as an advanced alternative to fixed-size token-based chunking (Kamradt, 2024). This approach is particularly effective for processing academic articles, where preserving the contextual integrity of information is essential. Traditional chunking methods, which divide text based on arbitrary token limits, often fragment meaningful sections of content, thereby disrupting the coherence and relevance of the retrieved material. By contrast, semantic chunking retains the relationships between sentences by grouping them according to their semantic similarity, ensuring that each chunk forms a semantically coherent unit. This is especially important when responding to user queries in an academic context, where the question might target a specific section of a paper, such as a hypothesis or conclusion, which must be retrieved in its entirety to maintain contextual accuracy. These findings are consistent with previous research suggesting that relationships between words within a sentence often convey more about the underlying semantic content than the individual words themselves (Ruas et al., 2020).

The process of semantic chunking begins by splitting the text into individual sentences, which are then converted into vector embeddings—representations of the semantic meaning of each sentence. These embeddings are compared using cosine similarity, and sentences with high similarity scores are grouped together into coherent chunks. A similarity threshold is crucial in this process, as it defines the boundaries between semantically distinct sections of text. When the similarity between consecutive sentences drops below this threshold, a new chunk is formed. This allows the system to dynamically determine where to place chunk boundaries based on the content itself, rather than relying on fixed token limits, thereby maintaining the logical flow of the text as originally intended by the authors.

In our study, it is also important to ensure that chunks are not only semantically coherent but also comparable in size compared to different experiments not using semantic chunking. Achieving a comparable word count across chunks was critical to ensure that the metrics used to evaluate the performance of our retrieval-augmented generation (RAG) system remained fair and directly comparable across different experimental conditions. To achieve this balance, we carefully adjusted the similarity threshold, ensuring that the resulting sentence clusters were of a similar word count, while still maintaining semantic coherence. This adjustment was key to ensuring that each chunk preserved its contextual relevance and that the retrieval process was optimized for both the relevance and coherence of the information. The selected threshold allowed us to balance the need for

coherent semantic groupings with the need for consistency in chunk size, ultimately contributing to the overall performance of the RAG system by enhancing both the precision of retrieval and the clarity of the generated responses.

3.2.4 Abstract-first Method

In the fourth stage of our pipeline, we implemented an abstract-first search strategy to enhance the effectiveness of the retrieval process. This approach takes advantage of the structured nature of academic articles, where the abstract offers a concise summary of the research findings, methods, and key conclusions. By focusing on abstracts, we can assess the relevance of a large number of articles quickly and more effectively before delving into full-text searches. Specifically, we constructed a vector database consisting solely of the abstracts of approximately 5,000 articles. The query is embedded and used to search within this abstract database, with the similarity between the query and the abstract embeddings calculated. From this search, the top 100 most relevant abstracts are selected for further investigation.

Once the top 100 abstracts are identified, we then leverage the titles of these articles to guide a more targeted search within the full-text vector database. Instead of searching through the entire body of 5,000 articles, the system focuses only on the titles of the top 100 most relevant documents, conducting a more focused retrieval process. This step uses the title metadata to locate relevant full-text chunks, allowing the system to evaluate detailed content only from the most promising articles, which have already been filtered based on their abstracts. This method provides a layered approach to retrieval: first narrowing down the search space using abstracts, and then refining it further using full-text analysis of a smaller subset of documents.

The abstract-first strategy significantly improves the precision of the system, as it avoids the potential distractions of unrelated full-text content by focusing only on the most semantically relevant documents. Additionally, by ensuring that only the most relevant 100 articles are selected from the initial search, we maintain a focused and relevant search scope when analyzing full-text content. As detailed in the experimental analysis we will discuss later, this approach helps avoid the pitfalls of excessive or irrelevant retrieval, contributing to more accurate and contextually meaningful outputs from the system while maintaining the necessary depth of analysis.

3.2.5 Enhanced Prompting Technique

In the final stage, we focused on optimizing the interaction between the Large Language Model (LLM) and the retrieved context by experimenting with various prompting techniques, informed by existing research on prompt engineering. The initial prompt provided by LangChain, which served as our baseline, was straightforward and task-oriented: "You are an assistant for question-answering tasks. Use the following pieces of retrieved-context to answer the question. If you don't know the answer, just say that you don't know. Use three sentences maximum and keep the answer concise." (Chase, 2022). While effective, this prompt was designed primarily for general question-answering tasks and lacked elements that could potentially further refine the quality of responses in more specialized scenarios. To address this, we explored more advanced prompting techniques that incorporate motivational and reward-based elements, building upon findings in the literature.

Specifically, the enhancement of our prompt was guided by studies that demonstrated the impact of emotional and incentive-based prompting on LLM performance. For example, Salinas and Morstatter (2024) showed that small modifications to the wording of prompts, such as including incentives or tip-offering, could significantly improve the precision and relevance of LLM-generated outputs. Additionally, emotional prompting techniques, as discussed by Li et al. (2023), provided further inspiration. Li's research indicates that prompts infused with positive reinforcement and emotional cues can elicit more contextually rich and nuanced responses from LLMs. While the emotional depth required for certain applications, such as mental health support, was not the primary focus of our system, the principle of encouraging more thoughtful and deliberate responses through positive reinforcement was highly relevant.

Inspired by these findings, we modified the original LangChain prompt to incorporate a hypothetical reward: "You are the best assistant for question-answering tasks. Your role is to answer the question excellently using the provided context. Use the following pieces of retrieved-context to answer the question. If you don't know the answer, just say that you don't know. Use three sentences maximum and keep the answer concise. I will tip you 1000 dollars for a perfect response." This variation was designed to encourage the LLM to approach the task with a higher degree of accuracy by framing the request as a challenge with a reward-based outcome. The final prompt was designed to assess the impact on the LLM's ability to generate responses that were both relevant to the user's query and faithful to the retrieved context.

3.3 Evaluation Using RAGAS

To assess the performance of our enhanced RAG application, we employed the RAGAS (Retrieval-Augmented Generation Assessment System) framework, which evaluates the quality of generated outputs based on three key metrics: Context Relevance, Faithfulness, and Answer Relevance (Es et al., 2023). Each of these metrics plays a critical role in ensuring that the generated responses are accurate, relevant, and grounded in the retrieved context. The RAGAS framework utilizes the OpenAI API to automatically determine metric values, such as identifying which sentences in the retrieved context are relevant to answering the given question or which claims in the generated answer can be inferred from the given context, based on predefined prompts outlined in the RAGAS paper. Since, unfortunately, there is no benchmark test set or established ground truth available for this kind of study, we had to create a custom test set and use the RAGAS framework, which relies on LLMs in the background as evaluators. Given these limitations, this approach was the only viable option to systematically assess the performance of our proposed architecture.

Context Relevance measures the relevance of the retrieved context, which is essential because the context should contain only the information needed to answer the provided query. The metric ranges from 0 to 1, with higher values indicating better relevance. To compute Context Relevance, we identified the sentences within the retrieved context that are pertinent to the query:

$$(1) \quad \text{Context Relevance} = \frac{|S|}{\text{Total number of sentences in retrieved context}}$$

Where $|S|$ is the number of sentences in the retrieved context that are relevant to answering the given question.

A high Context Relevance score indicates that the *retrieved context* is highly focused and contains minimal extraneous information. This metric is particularly important in our application because our enhancements aim to optimize the retrieval and presentation of contextually accurate information.

Faithfulness assesses the factual consistency of the *generated answer* with the *retrieved context*. It is calculated by comparing the claims made in the generated answer with

the information provided in the context. The metric also ranges from 0 to 1, with higher scores indicating better factual accuracy. This metric ensures that the generated response is grounded in the retrieved context, avoiding hallucinations and maintaining factual accuracy. For our application, Faithfulness is also crucial because it directly impacts the reliability of the responses provided to data scientists.

$$(2) \quad \text{Faithfulness score} = \frac{\text{Number of claims in the generated answer that can be inferred from the given context}}{\text{Total number of claims in the generated answer}}$$

Answer Relevance evaluates how well the *generated answer* addresses the *provided query*. This metric is assessed by computing the cosine similarity between the embeddings of the original question and the generated question, which is reverse-engineered from the answer. A higher Answer Relevance score indicates that the generated answer closely aligns with the original question, ensuring that the response is relevant and directly addresses the user’s query. In our context, Answer Relevance is vital for assessing the overall effectiveness of the RAG application in producing useful and accurate responses.

$$(3) \quad \text{Answer Relevance} = \frac{1}{N} \sum_{i=1}^N \cos(E_{g_i}, E_o)$$

Where:

- E_{g_i} is the embedding of the generated question i .
- E_o is the embedding of the original question.
- N is the number of generated questions

Using these detailed metrics, the RAGAS framework provides a comprehensive evaluation of the performance of our enhanced RAG system. Each metric measures performance in different but complementary areas. Figure 3.3 visually illustrates the aspects that these metrics cover. The metrics not only validate the improvements we have made but also offer insights into areas where further enhancements can be pursued, ensuring that the application continues to deliver high-quality, reliable, and contextually relevant answers to data scientists.

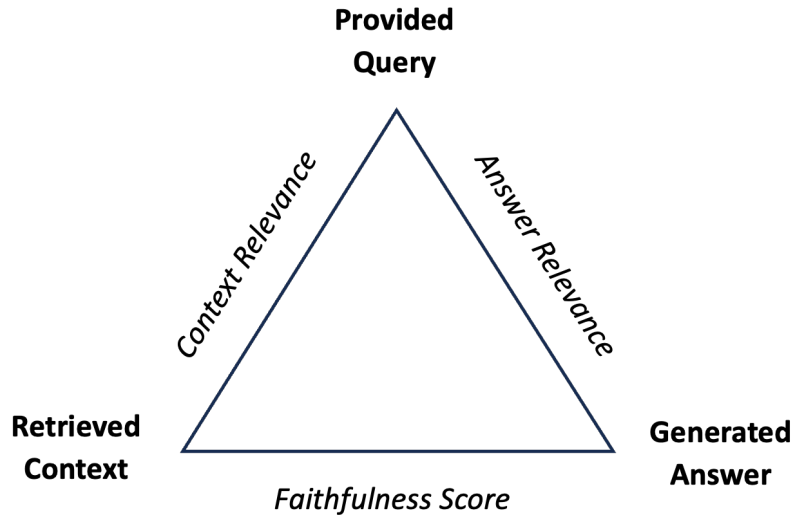


Figure 3.3 Three Performance Metrics of RAGAS Framework and Their Corresponding Aspects

We have also incorporated the average word count into the analysis and displayed it in the results table to consider its impact on these metrics. Specifically, metrics like context relevance are highly dependent on the number of words in the retrieved contexts. When more sentences are included in a context, it becomes more challenging to achieve high scores across these metrics because, for context relevance to be maximized, all sentences should be relevant to the question. With more sentences, the likelihood of including unrelated content increases, potentially lowering the overall score.

4. RESULTS

The results of our experiments, summarized in the tables, demonstrate the impact of various enhancements to the RAG pipeline on key metrics: Context Relevance (CR), Faithfulness, and Answer Relevance. Each metric’s value is averaged over 1,500 replications -30 replications for each one of the 50 questions used in the analysis. The choice of 30 replications per question is grounded in the Central Limit Theorem (CLT), which asserts that with a sample size of 30, the distribution of the sample means will approximate a normal distribution, ensuring the robustness of our results (Moore and McCabe, 1989).

4.1 Performance Metrics for Different Fine-tuning Experiments

Table 4.1 Performance Metrics for Different Fine-tuning Experiments. (FT = Fine-tuning, TB = Textbooks, G = GROBID, SNS = Semantic Node Splitter)

Experiment	Context Relevance	Faithfulness	Answer Relevance	Average Word Count
1. No FT	0.302	0.687	0.888	853.88
2. 5 TB	0.337	0.658	0.871	853.90
3. 17 TB	0.372	0.756	0.902	836.10
4. 5 TB + G	0.062	0.198	0.790	913.00
5. 17 TB + G	0.335	0.740	0.909	839.34
6. 17 TB + G + SNS	0.339	0.701	0.902	841.72

In Table 4.1, all experiments, including *No FT* case which corresponds to $B + G$ in Table 4.3), the models incorporated GROBID for preprocessing and fine-tuning to stan-

standardize data structuring. Variations in performance were solely driven by differences in the fine-tuning data, particularly the amount and type of training materials. Without fine-tuning, the baseline model incorporated with GROBID exhibited relatively low Context Relevance (CR) at 0.0302, despite relatively high Faithfulness (0.687) and Answer Relevance (0.888). This highlighted the necessity of fine-tuning to improve the model’s ability to retrieve relevant content.

Fine-tuning on 5 textbooks (Experiment 2) significantly improved CR to 0.337, with a slight rise in Faithfulness and a small decline in Answer Relevance. When expanded to 17 textbooks (Experiment 3), CR reached 0.372, demonstrating the advantage of a larger training dataset, with improved Faithfulness (0.756) and Answer Relevance (0.902). This model was selected for further testing due to its strong overall performance in CR.

In Experiment 4, adding GROBID to fine-tune with 5 textbooks caused a drop in CR to 0.062, likely due to the removal of important content during preprocessing. However, combining GROBID with fine-tuning with 17 textbooks (Experiment 5) restored CR to 0.335, improving both Faithfulness and Answer Relevance. Further improvements were seen in Experiment 6, where the addition of Semantic Node Splitter (SNS) slightly increased CR to 0.339, maintaining competitive Faithfulness and high Answer Relevance.

Overall, fine-tuning with larger datasets significantly improved CR, while GROBID and SNS added further refinements. The models from Experiments 3 (17 TB) and 6 (17 TB + G + SNS) were chosen for further evaluation, having demonstrated the best performance in retrieving relevant academic context, which is the main focus of this study.

Table 4.2 Significant Pairwise Comparisons for Context Relevance Using Tukey’s HSD Test.

Comparison	Mean Difference	p-value	Significant?
1 vs. 2	0.035	<0.001	Yes
2 vs. 3	0.035	<0.001	Yes
3 vs. 4	-0.310	<0.001	Yes
3 vs. 5	-0.037	<0.001	Yes
3 vs. 6	-0.034	<0.001	Yes
6 vs. 2	-0.002	0.999	No
6 vs. 5	-0.003	0.994	No
2 vs. 5	-0.002	0.9996	No

Context Relevance was utilized exclusively to perform the significance tests because the

primary aim is to assess the impact of fine-tuning on the model’s ability to retrieve relevant content and adapt the embedding model effectively. Since it is crucial for the embedding to accurately understand the query and obtain the relevant chunks, focusing on Context Relevance allows us to detect the contributions of different fine-tuning experiments more precisely. To statistically evaluate the significance of the differences in Context Relevance across configurations, we employed Welch’s ANOVA followed by Tukey’s HSD test for post-hoc analysis (Welch, 1951) (Tukey, 1949). Welch’s ANOVA was chosen because it does not assume equal variances, making it a more robust method when the homogeneity of variances assumption is violated. This approach ensures the reliability of the results by accounting for variance differences across the experiments. Tukey’s HSD test was then applied to identify which specific pairwise comparisons showed statistically significant differences, providing a detailed understanding of the performance variations between the configurations.

The results of Tukey’s HSD tests presented in Table 4.2 indicate a statistically significant (with $p < 0.001$) enhancement in Context Relevance from the baseline Experiment 1 (No FT) to Experiment 3 (17 TB). Evidently, Experiment 3 exhibits a statistically significant superiority over all other experimental conditions with respect to Context Relevance. While the results for Experiment 6 (17 TB + G + SNS), Experiment 2 (5 TB), and Experiment 5 (17 TB + G) are closely aligned, exhibiting no statistically significant differences among them as detailed in the last three rows of Table 4.2, we have chosen to further evaluate Experiment 6 alongside Experiment 3 due to its position as the second highest in overall average Context Relevance.

4.2 Performance Metrics for RAG Pipeline Configurations Using the

Fine-tuning Model with Experiment 3

The performance metrics for different RAG pipeline configurations using the fine-tuned model from Experiment 3, as shown in Table 4.3, illustrate how various enhancements affect the retrieval process. Adding GROBID (B + G) significantly improved Context Relevance (CR) from 0.031 to 0.302, demonstrating its effectiveness in structuring the data. Faithfulness increased to 0.687, while Answer Relevance slightly declined to 0.888, indicating some variability in the use of the retrieved context. The average word count also rose to 853.88, reflecting broader context retrieval with GROBID.

Table 4.3 Performance Metrics for Different RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 3. (B = Baseline, G = GROBID, FT = Fine-tuning, SC = Semantic Chunking, AF = Abstract First, EPT = Enhanced Prompting Technique)

Configuration	Context Relevance	Faithfulness	Answer Relevance	Average Word Count
1. B	0.031	0.622	0.898	772.92
2. B + G	0.302	0.687	0.888	853.88
3. B + G + FT	0.372	0.756	0.902	836.10
4. B + G + FT + SC	0.456	0.599	0.852	947.94
5. B + G + FT + SC + AF	0.447	0.606	0.848	915.80
6. B + G + FT + SC + AF + EPT	0.454	0.588	0.887	917.86

Combining GROBID with fine-tuning (B + G + FT) further improved CR to 0.372, with higher Faithfulness (0.756) and Answer Relevance (0.902). This configuration had the highest Faithfulness, suggesting efficient use of the relevant context. The average word count dropped to 836.10, reflecting a more refined retrieval process.

Semantic Chunking (B + G + FT + SC) raised CR to 0.456, improving retrieval of semantically coherent chunks, though Faithfulness dropped to 0.599, likely due to the larger average word count (947.94). Adding the Abstract-First strategy (B + G + FT + SC + AF) slightly reduced CR to 0.447 but improved Faithfulness to 0.606, with a lower average word count of 915.80, indicating more concise retrieval.

The final configuration (B + G + FT + SC + AF + EPT) maintained a high CR (0.454) and improved Answer Relevance to 0.887. Faithfulness slightly decreased to 0.588, suggesting the potential for better alignment between context and generated responses. Overall, the fine-tuning model from Experiment 2 performed best when GROBID, fine-tuning, and Semantic Chunking were combined, while Enhanced Prompting improved the quality of responses.

The significant pairwise comparisons presented in Table 4.4 illustrate the impact of each RAG pipeline enhancement on the key performance metrics. Until the fourth configuration, the successive improvements in Context Relevance and Faithfulness were statistically significant, as evidenced by the p-values below 0.001 for each step. This indicates that adding GROBID, fine-tuning, and Semantic Chunking sequentially contributed meaningfully to the improvement in these metrics. However, the Abstract-First strategy and Enhanced Prompting Technique (configurations 5 and 6) did not produce statistically significant changes in Context Relevance or Faithfulness, suggesting that these strategies may not further optimize these aspects. In contrast, Answer Relevance only showed significant differences in configurations 4 and 6, corresponding to the application of Se-

Table 4.4 Significant Pairwise Comparisons for Different Metrics Using Tukey’s HSD Test for Table 4.3. (CR = Context Relevance, F = Faithfulness, AR = Answer Relevance)

Metric	Comparison	Mean Difference	p-value	Significant?
Context Relevance	1 vs. 2	0.2713	<0.001	Yes
	2 vs. 3	0.0704	<0.001	Yes
	3 vs. 4	0.0838	<0.001	Yes
	4 vs. 5	-0.0089	0.952	No
	5 vs. 6	0.0073	0.980	No
Faithfulness	1 vs. 2	0.0651	<0.001	Yes
	2 vs. 3	0.0691	<0.001	Yes
	3 vs. 4	-0.1576	<0.001	Yes
	4 vs. 5	0.0068	0.992	No
	5 vs. 6	-0.0171	0.673	No
Answer Relevance	1 vs. 2	-0.0103	0.454	No
	2 vs. 3	0.0141	0.126	No
	3 vs. 4	-0.0496	<0.001	Yes
	4 vs. 5	-0.0036	0.989	No
	5 vs. 6	0.0380	<0.001	Yes

mantic Chunking and the Enhanced Prompting Technique. This is consistent with the role of Semantic Chunking in retrieving more relevant context, thereby improving the relevance of the generated responses, and the Enhanced Prompting Technique’s ability to refine the model’s output, yielding more precise answers to the posed questions.

4.3 Performance Metrics for RAG Pipeline Configurations Using the

Fine-tuning Model with Experiment 6

Table 4.5 Performance Metrics for Different RAG Pipeline Configurations Using the Fine-tuning Model with Experiment 6. (B = Baseline, G = GROBID, FT = Fine-tuning, SC = Semantic Chunking, AF = Abstract First, EPT = Enhanced Prompting Technique)

Configuration	Context Relevance	Faithfulness	Answer Relevance	Average Word Count
1. B	0.031	0.622	0.898	772.92
2. B + G	0.302	0.687	0.888	853.88
3. B + G + FT	0.339	0.701	0.902	841.72
4. B + G + FT + SC	0.475	0.618	0.872	840.24
5. B + G + FT + SC + AF	0.507	0.592	0.863	753.16
6. B + G + FT + SC + AF + EPT	0.507	0.592	0.890	753.16

The results in Table 4.5 present the performance of various RAG pipeline configurations using the fine-tuned model from Experiment 6 across key metrics. The baseline (B) configuration shows low Context Relevance (CR) at 0.031, with higher Faithfulness (0.622) and Answer Relevance (0.898), but the retrieved context remains minimally effective. Adding GROBID (B + G) improves CR to 0.302, with a slight increase in Faithfulness (0.687) and a small decrease in Answer Relevance (0.888). The average word count also rises to 853.88, reflecting broader content retrieval.

Introducing fine-tuning alongside GROBID (B + G + FT) further increases CR to 0.339, with Faithfulness rising to 0.701 and Answer Relevance stabilizing at 0.902. Semantic Chunking (B + G + FT + SC) brings the highest CR yet (0.485), though Faithfulness drops slightly to 0.640, with Answer Relevance decreasing to 0.874. The Abstract-First strategy (B + G + FT + SC + AF) improves CR to 0.507, though both Faithfulness (0.592) and Answer Relevance (0.863) decrease, possibly due to the narrower focus on abstract-level content.

Finally, integrating the Enhanced Prompting Technique (B + G + FT + SC + AF + EPT) maintains CR at 0.507 while increasing Answer Relevance to 0.890. This suggests that enhanced prompts help the model generate more coherent responses, even as Faithfulness remains stable at 0.592. Overall, the best performance in CR was achieved with

Semantic Chunking and Abstract-First strategies, while enhanced prompting improved answer quality.

Table 4.6 Significant Pairwise Comparisons for Different Metrics Using Tukey’s HSD Test for Table 4.5. (CR = Context Relevance, F = Faithfulness, AR = Answer Relevance)

Metric	Comparison	Mean Difference	p-value	Significant?
Context Relevance	1 vs. 2	0.2713	<0.001	Yes
	2 vs. 3	0.0367	0.0065	Yes
	3 vs. 4	0.1374	<0.001	Yes
	4 vs. 5	0.0315	0.0342	Yes
	5 vs. 6	-0.0002	1.0000	No
Faithfulness	1 vs. 2	0.0651	<0.001	Yes
	2 vs. 3	0.0137	0.834	No
	3 vs. 4	-0.0868	<0.001	Yes
	4 vs. 5	-0.0224	0.3643	No
	5 vs. 6	0.0002	1.0000	No
Answer Relevance	1 vs. 2	-0.0103	0.3491	No
	2 vs. 3	0.0140	0.0736	No
	3 vs. 4	-0.0295	<0.001	Yes
	4 vs. 5	-0.0090	0.5050	No
	5 vs. 6	0.0270	<0.001	Yes

The significant pairwise comparisons in Table 4.6 indicate consistent improvements in Context Relevance across each step, up until the final configuration involving Enhanced Prompting. This aligns with expectations, as Enhanced Prompting primarily influences the response generation phase rather than context retrieval. The Abstract-First approach (configuration 5) produced a significant enhancement in Context Relevance, contrasting with its lack of significant effect in Experiment 3 (Table 4.4). This suggests that Abstract-First may facilitate more efficient retrieval processes in certain experimental conditions. Faithfulness significantly increased only with the initial addition of GROBID (configuration 2), while it showed a notable decline with the introduction of Semantic Chunking (configuration 4), possibly due to the increased length of retrieved content impacting alignment with the generated response. For Answer Relevance, the metric significantly decreased with the addition of Semantic Chunking but later improved with Enhanced Prompting, reflecting the latter’s role in refining the generated responses to align more closely with the retrieved context.

5. DISCUSSION & CONCLUSIONS

The experiments conducted in this study aimed to refine the Retrieval-Augmented Generation (RAG) pipeline for academic literature, focusing on enhancing the retrieval of context relevant to data science-related queries. Several configurations were evaluated across key metrics, including Context Relevance (CR), Faithfulness, and Answer Relevance, providing insight into how different components of the pipeline affect its overall performance. The emphasis of this work was on improving CR, given its direct correlation with the ability to retrieve pertinent context from large corpora of academic literature.

One of the key findings is the positive impact of fine-tuning on CR across all configurations. Models that were fine-tuned using domain-specific textbooks, particularly in Experiments 2 and 5, consistently outperformed the baseline model in retrieving context relevant to the query. This highlights the importance of curating high-quality, domain-specific datasets when fine-tuning language models to optimize their performance for specialized tasks. The results also suggest that increasing the size of the fine-tuning corpus, as seen in Experiment 2, improves not only CR but also Faithfulness and Answer Relevance. This trend underscores the value of larger and more comprehensive training datasets for improving the system’s ability to retrieve and use relevant information.

The introduction of GROBID provided mixed results. While it significantly improved CR by structuring the data for more precise retrieval, its incorporation without sufficient fine-tuning, as seen in Experiment 3, led to a noticeable drop in CR. This suggests that while GROBID’s document structuring is beneficial, it must be paired with ample training data to ensure that critical context is not inadvertently removed during preprocessing. The combination of GROBID with fine-tuning on larger datasets restored CR to higher levels, as observed in Experiments 4 and 5, indicating that the issues introduced by GROBID can be mitigated with appropriate data preparation (Lopez, 2024).

Semantic Chunking also played a significant role in improving CR, particularly in the final stages of the pipeline. By grouping text into semantically coherent chunks, this technique ensured that the model retrieved more meaningful and contextually relevant

information. However, the trade-off between longer retrieved contexts and Faithfulness suggests that future work could explore ways to balance chunk size with the system’s ability to generate precise answers from the retrieved information. This is evident in the slight reduction in Faithfulness observed when Semantic Chunking was introduced.

The Abstract-First strategy introduced in the pipeline further refined the retrieval process by focusing on abstracts to filter relevant content before conducting a more in-depth full-text search. This approach led to notable improvements in CR, achieving the highest scores across all configurations. Interestingly, the use of abstracts challenges the conventional assumption that broader searches yield better results. The findings suggest that a more focused search using abstract-level information can be equally, if not more, effective than full-text searches. This method not only reduced the average word count but also demonstrated that a concise, abstract-based search can still maintain high relevance in context retrieval. However, it remains to be seen whether this strategy is as effective in other domains where abstracts may not fully capture the nuances of the article’s content, presenting an area for future research.

The Enhanced Prompting Technique, while not affecting CR directly, improved Answer Relevance by guiding the LLM to better utilize the retrieved context when generating answers. Since the LLM combines the retrieved context with the query to produce the response, the prompt’s structure plays a pivotal role in maximizing the utility of the language model. The relatively stable CR across configurations that employed enhanced prompting confirms that this technique’s impact is not on retrieval but rather on improving how well the model aligns its responses with the given context. This result underscores the importance of prompt design in optimizing response generation. Moreover, there remains room for further improvements in prompt alignment, as the open-ended nature of prompts offers opportunities to refine how context and query are combined to generate more precise and contextually accurate responses.

In future enhancements, incorporating knowledge graph-based approaches for automating the validation and structuring of retrieved content could significantly improve the accuracy and reliability of RAG applications. Such techniques, as demonstrated in the construction and validation of behavioral models using automated knowledge extraction, show promise for dynamically synthesizing robust and relevant information from diverse sources (Sonnenschein et al., 2024). By integrating these methods, RAG systems could better manage complex information retrieval tasks, thereby further enhancing decision-making processes across various domains.

In conclusion, this study illustrates that the proposed Enhanced RAG Architecture, which

integrates fine-tuning on domain-specific datasets, employs GROBID for data structuring, implements Semantic Chunking, and utilizes an Abstract-First strategy, significantly enhances the capability of the RAG pipeline to retrieve pertinent academic content. The fine-tuning models with larger and more diverse training sets, as demonstrated in Experiments 3 and 6, were particularly effective in increasing Context Relevance (CR), which aligns with the study’s goal of improving retrieval of academic literature for data scientists. While Semantic Chunking and Abstract-First approaches further refined retrieval precision, the observed trade-offs in Faithfulness suggest a need for future work to optimize how retrieved content is utilized in generating responses. The Enhanced Prompting Technique contributed to more contextually aligned answers, highlighting the role of prompt design in enhancing LLM performance within RAG systems. These findings offer useful insights for advancing context retrieval systems in academic research, with further investigations needed to assess the broader applicability of these methods.

Despite its contributions, this study has several limitations. One key challenge is the lack of a standardized ground truth for evaluating the quality of the generated answers. To address this, we utilized the RAGAS framework, which relies on LLMs for evaluation. While this method was the most viable given the circumstances, it introduces a degree of subjectivity into the assessment, potentially affecting the interpretation of the results. Additionally, the reliance on a custom test set, developed specifically for this study, poses another limitation. Although the questions were designed to reflect realistic challenges in data science, the lack of a widely accepted benchmark for RAG evaluation in academic literature restricts the generalizability of the findings. Future research should aim to establish standardized benchmarks and ground truth datasets to facilitate more robust and comparable evaluations across studies.

BIBLIOGRAPHY

- Agarwal, S., Laradji, I. H., Charlin, L., and Pal, C. (2024). Litllm: A toolkit for scientific literature review. *arXiv*. Retrieved from <https://arxiv.org/abs/2402.01788>.
- Alshammari, S., Basalelah, L., Rukbah, W. A., Alsuhibani, A., and Wijesinghe, D. S. (2023). Knimezobot: Enhancing literature review with zotero and knime openai integration using retrieval-augmented generation. *arXiv*. Retrieved from <https://arxiv.org/abs/2311.04310>.
- Arora, D., Kini, A., Chowdhury, S. R., Natarajan, N., Sinha, G., and Sharma, A. (2023). Gar-meets-rag paradigm for zero-shot information retrieval. *arXiv*. Retrieved from <https://arxiv.org/abs/2310.20158>.
- Boyd, J. (2015). Automatic metadata extraction - the high energy physics use case. Master’s thesis, Ecole Polytechnique, Lausanne. Retrieved from <https://cds.cern.ch/record/2039361.Presented04Sep2015>.
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., and Wirth, R. (2000). Crisp-dm 1.0 step-by-step data mining guide (august 2000). *The CRISP-DM consortium*. Retrieved from <http://www.crisp-dm.org/CRISPWP-0800.pdf>.
- Chase, H. (2022). Langchain. *GitHub Repository*. Retrieved from <https://github.com/hwchase17/langchain>.
- Dušek, R., Galias, C., Wojciechowska, L., and Wawer, A. (2023). Improving domain-specific retrieval by nli fine-tuning. In *2023 18th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pages 949–953. IEEE. Retrieved from <https://doi.org/10.15439/2023F7569>.
- Es, S., James, J., Espinosa-Anke, L., and Schockaert, S. (2023). Ragas: Automated evaluation of retrieval augmented generation. *arXiv*. Retrieved from <https://arxiv.org/abs/2309.15217>.
- Gao, T., Yen, H., Yu, J., and Chen, D. (2023). Enabling large language models to generate text with citations. *arXiv*. Retrieved from <https://arxiv.org/abs/2305.14627>.
- Gunasekar, S., Zhang, Y., Aneja, J., Mendes, C. C. T., Del Giorno, A., Gopi, S., and Li, Y. (2023). Textbooks are all you need. *arXiv*. Retrieved from <https://arxiv.org/abs/2306.11644>.
- Han, W., Shen, J., Liu, Y., Shi, Z., Xu, J., Hu, F., and Ge, M. (2024). Legalsst: Human-centered and ai-empowered machine to enhance court productivity and legal assistance. *Information Sciences*, 679:121052. Retrieved from <https://doi.org/10.1016/j.ins.2024.121052>.

- Izacard, G. and Grave, E. (2020). Leveraging passage retrieval with generative models for open domain question answering. *arXiv*. Retrieved from <https://arxiv.org/abs/2007.01282>.
- Kaesberg, L. B., Ruas, T., Wahle, J. P., and Gipp, B. (2024). Citeassist: A system for automated preprint citation and bibtex generation. *arXiv*. Retrieved from <https://arxiv.org/abs/2407.03192>.
- Kamradt, G. (2024). Semantic chunking. *FullStack Retrieval Tutorials*. Retrieved from <https://github.com/FullStackRetrieval-com/RetrievalTutorials/tree/main/tutorials/LevelsOfTextSplitting>.
- Khouya, N., Retbi, A., and Bennani, S. (2024). Contextual revelation of intentions in machine learning articles on arxiv: A new dimension in academic content analysis. In *2024 IEEE 12th International Symposium on Signal, Image, Video and Communications (ISIVC)*, pages 1–4. IEEE.
- Krenn, M., Buffoni, L., Coutinho, B., Eppel, S., Foster, J. G., Gritsevskiy, A., and Kopp, M. (2022). Predicting the future of ai with ai: High-quality link prediction in an exponentially growing knowledge network. *arXiv*. Retrieved from <https://arxiv.org/abs/2210.00881>.
- Levonian, Z., Li, C., Zhu, W., Gade, A., Henkel, O., Postle, M. E., and Xing, W. (2023). Retrieval-augmented generation to improve math question-answering: Trade-offs between groundedness and human preference. *arXiv*. Retrieved from <https://arxiv.org/abs/2310.03184>.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., and Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474.
- Li, C., Wang, J., Zhang, Y., Zhu, K., Hou, W., Lian, J., and Xie, X. (2023). Large language models understand and can be enhanced by emotional stimuli. *arXiv*. Retrieved from <https://arxiv.org/abs/2307.11760>.
- Li, Y., Zhao, J., Li, M., Dang, Y., Yu, E., Li, J., and Tao, C. (2024). Refai: a gpt-powered retrieval-augmented generative tool for biomedical literature recommendation and summarization. *Journal of the American Medical Informatics Association*, ocae129. Retrieved from <https://doi.org/10.1093/jamia/ocae129>.
- Liu, J. (2022). Llamaindex: Building a gpt-powered semantic search index from your data. *GitHub Repository*. Retrieved from https://github.com/jerryjliu/llama_index.
- Lopez, P. (2008-2024). Grobid. *GitHub Repository*. Retrieved from <https://github.com/kermitt2/grobid>.
- Moore, D. S. and McCabe, G. P. (1989). *Introduction to the practice of statistics*. WH Freeman/Times Books/Henry Holt & Co.

- Mysore, S., Jasim, M., Song, H., Akbar, S., Randall, A. K. C., and Mahyar, N. (2023). How data scientists review the scholarly literature. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*, pages 137–152.
- Nguyen, Z., Annunziata, A., Luong, V., Dinh, S., Le, Q., Ha, A. H., and Nguyen, C. (2024). Enhancing q&a with domain-specific fine-tuning and iterative reasoning: A comparative study. *arXiv*. Retrieved from <https://arxiv.org/abs/2404.11792>.
- Otto, W., Zloch, M., Gan, L., Karmakar, S., and Dietze, S. (2023). Gsap-ner: A novel task, corpus, and baseline for scholarly entity extraction focused on machine learning models and datasets. *arXiv*. Retrieved from <https://arxiv.org/abs/2311.09860>.
- Pu, Y., He, Z., Qiu, T., Wu, H., and Yu, B. (2024). Customized retrieval augmented generation and benchmarking for eda tool documentation qa. *arXiv*. Retrieved from <https://arxiv.org/abs/2407.15353>.
- Ruas, T., Ferreira, C. H. P., Grosky, W., de França, F. O., and de Medeiros, D. M. R. (2020). Enhanced word embeddings using multi-semantic representation through lexical chains. *Information Sciences*, 532:16–32. Retrieved from <https://doi.org/10.1016/j.ins.2020.04.048>.
- Salinas, A. and Morstatter, F. (2024). The butterfly effect of altering prompts: How small changes and jailbreaks affect large language model performance. *arXiv*. Retrieved from <https://arxiv.org/abs/2401.03729>.
- Sambar, M., Vázquez, G. R., Vázquez, A. V., and Vázquez, F. X. (2024). A chatgpt assisted reading protocol for undergraduate research students engaging with biophysics literature. *bioRxiv*. Retrieved from <https://doi.org/10.1101/2024.09.11.612473>.
- Shi, F., Chen, X., Misra, K., Scales, N., Dohan, D., Chi, E. H., and Zhou, D. (2023a). Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR.
- Shi, Y., Xu, S., Liu, Z., Liu, T., Li, X., and Liu, N. (2023b). Mededit: Model editing for medical question answering with external knowledge bases. *arXiv*. Retrieved from <https://arxiv.org/abs/2309.16035>.
- Siriwardhana, S., Weerasekera, R., Wen, E., Kaluarachchi, T., Rana, R., and Nanayakkara, S. (2023). Improving the domain adaptation of retrieval augmented generation (rag) models for open domain question answering. *Transactions of the Association for Computational Linguistics*, 11:1–17.
- Sonnenschein, T. S., de Wit, G. A., den Braver, N. R., Vermeulen, R. C., and Scheider, S. (2024). Validating and constructing behavioral models for simulation and projection using automated knowledge extraction. *Information Sciences*, 662:120232. Retrieved from <https://doi.org/10.1016/j.ins.2024.120232>.
- Tkaczyk, D., Collins, A., Sheridan, P., and Beel, J. (2018). Evaluation and comparison of open source bibliographic reference parsers: A business use case. *arXiv*. Retrieved from <https://arxiv.org/abs/1802.01168>.

- Tkaczyk, D., Szostek, P., Fedoryszak, M., Dendek, P. J., and Bolikowski, L. (2015). Cermine: automatic extraction of structured metadata from scientific literature. *International Journal on Document Analysis and Recognition (IJDAR)*, 18(4):317–335.
- Transformers, S. (2021). all-mpnet-base-v2 [pre-trained model]. Hugging Face. Retrieved from <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>.
- Tukey, J. W. (1949). Comparing individual means in the analysis of variance. *Biometrics*, 5(2):99–114. Retrieved from <https://doi.org/10.2307/3001913>.
- Welch, B. L. (1951). On the comparison of several mean values: An alternative approach. *Biometrika*, 38(3/4):330–336. Retrieved from <https://doi.org/10.2307/2332579>.
- Wiratunga, N., Abeyratne, R., Jayawardena, L., Martin, K., Massie, S., Nkisi-Orji, I., and Fleisch, B. (2024). Cbr-rag: Case-based reasoning for retrieval augmented generation in llms for legal question answering. In *International Conference on Case-Based Reasoning*, pages 445–460. Cham: Springer Nature Switzerland.
- Wong, E. Y. (2017). e-print archive: arxiv.org. *Technical Services Quarterly*, 34(1):111–113.
- Wu, J. and Mooney, R. J. (2022). Breaking down questions for outside-knowledge vqa. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Wu, J., Zhu, J., and Qi, Y. (2024). Medical graph rag: Towards safe medical large language model via graph retrieval-augmented generation. *arXiv*. Retrieved from <https://arxiv.org/abs/2408.04187>.

APPENDIX

A. Advanced-Data Science Questions

1. I'm a data scientist in the financial sector, selecting features to predict stock prices. I've used traditional methods like backward elimination. Which advanced feature selection technique should I consider to enhance model performance?
2. In the context of agriculture, I'm predicting crop yields based on soil and weather data. Having tried manual feature selection, what automated feature selection method would help in identifying the most influential factors for yield prediction?
3. As a data scientist in the insurance industry, I'm working on customer churn prediction. I've used correlation-based feature selection but need to account for nonlinear relationships. Which technique would be suitable for this purpose?
4. Working in the healthcare sector, I deal with patient vital signs data that often contains outliers. After using Z-score for outlier detection, what robust outlier elimination technique should I adopt for cleaner datasets?
5. In the manufacturing industry, I'm analyzing sensor data from machinery for predictive maintenance. Having used the IQR method to detect outliers, which advanced technique would you recommend for more accurate anomaly detection and elimination?
6. I'm analyzing customer behavior data in e-commerce, using PCA for dimensionality reduction. Can you suggest a modern technique that might offer better insights while preserving data variance?
7. In environmental science, I deal with high-dimensional climate data. Having used t-SNE for visualization, what other dimensionality reduction method should I consider for more effective data analysis?

8. As a data scientist in the telecom sector, I work with extensive network traffic data. After applying LDA for feature extraction, what newer technique would enhance the interpretability and accuracy of my models?
9. In the energy industry, I'm analyzing sensor data from smart grids. While PCA has helped in reducing dimensions, what alternative should I explore for better performance with large datasets?
10. I'm studying brain imaging data for neurological research. Having tried factor analysis, what other dimensionality reduction technique could help in uncovering hidden patterns more effectively?
11. In the automotive industry, I develop models for autonomous vehicles using real-world driving data. What synthetic data generation technique can I use to augment my dataset while ensuring realistic scenarios?
12. Working with confidential financial data, I need to generate synthetic datasets for model training. Which method would provide a balance between data utility and privacy?
13. As a data scientist in social sciences, I deal with sensitive survey data. What technique should I use to generate synthetic data that maintains the statistical properties of the original data?
14. In robotics, I'm training a robot to navigate complex environments. Having used Q-learning, what advanced reinforcement learning algorithm should I explore for improved performance in dynamic settings?
15. As a game developer, I'm using reinforcement learning to create intelligent agents. Beyond DQN, which algorithm would enhance the adaptability and learning speed of my agents?
16. In the pharmaceutical industry, I'm classifying drug responses based on genetic data. After using SVM, what advanced classification algorithm should I consider for better accuracy and interpretability?
17. As a data scientist in the automotive sector, I'm classifying vehicle types from sensor data. Having used decision trees, which more sophisticated method could improve classification performance?
18. Working in cybersecurity, I classify network traffic for threat detection. Beyond random forests, what advanced classification technique should I employ for more

accurate threat identification?

19. In the education sector, I'm classifying student performance based on learning behaviors. After using k-NN, which algorithm would offer better results considering interpretability and accuracy?
20. As a data scientist in the legal domain, I'm classifying legal documents by type. What advanced classification method would enhance my document categorization system?
21. In the energy sector, I'm predicting energy consumption using multiple regression. What advanced regression model could provide better accuracy and handle nonlinearity more effectively?
22. Working in the telecommunications industry, I'm predicting customer call durations. Having tried ridge regression, which sophisticated regression technique should I explore next?
23. In the field of epidemiology, I'm predicting disease spread based on environmental factors. Beyond polynomial regression, what model should I use to improve predictive accuracy?
24. As a data scientist in real estate, I'm predicting property values. After using lasso regression, which advanced method would you recommend for incorporating spatial dependencies?
25. In transportation, I'm predicting traffic flow using regression models. What advanced technique should I consider for capturing complex patterns in traffic data?
26. In the context of smart cities, I'm clustering IoT sensor data. Having used DBSCAN, what other clustering algorithm would be effective for large, diverse datasets?
27. As a data scientist in the food industry, I'm clustering consumer taste preferences. Beyond hierarchical clustering, which algorithm should I explore for better segmentation?
28. Working in the entertainment sector, I'm clustering movie preferences. What advanced clustering method can help in identifying more nuanced audience segments?
29. In environmental monitoring, I'm clustering air quality data. After using agglomerative clustering, which technique would provide better results for large-scale environmental datasets?

30. As a data scientist in sports analytics, I'm clustering player performance metrics. Beyond Gaussian Mixture Models, which advanced clustering method should I consider for improved insights?
31. In medical imaging, I'm detecting abnormalities in MRI scans. Beyond CNNs, what advanced deep learning architecture should I explore for higher accuracy?
32. As a data scientist in agriculture, I'm analyzing drone imagery for crop health. Having used image segmentation, which technique should I adopt for more precise analysis?
33. Working with satellite imagery in environmental science, I'm classifying land use. Beyond traditional CNNs, which model should I explore for improved accuracy and efficiency?
34. In the automotive sector, I'm developing object detection systems for autonomous vehicles. What advanced image processing technique would enhance detection accuracy in real-time scenarios?
35. As a data scientist in fashion, I'm analyzing images for style recognition. Beyond transfer learning, which approach should I use to improve model performance?
36. In legal tech, I'm analyzing contracts using NLP. Having used word embeddings, what advanced technique should I consider for better understanding legal jargon and nuances?
37. As a data scientist in customer service, I'm developing chatbots for automated responses. Beyond basic sequence models, which advanced NLP model should I explore for more natural interactions?
38. Working in publishing, I'm analyzing book reviews for sentiment analysis. What state-of-the-art NLP technique would help in capturing complex sentiments more accurately?
39. In the field of education, I'm developing automated grading systems. Having used traditional NLP methods, which advanced model should I adopt for better accuracy in grading essays?
40. As a data scientist in social media analysis, I'm detecting trends from tweets. Beyond LSTM, which cutting-edge NLP technique should I consider for more insightful trend analysis?

41. In finance, I'm forecasting stock prices using traditional ARIMA models. What advanced time series analysis method should I use to better capture market volatility?
42. As a data scientist in retail, I'm predicting sales trends. Beyond SARIMA, which technique should I explore to account for seasonal variations and promotions?
43. Working in the energy sector, I'm forecasting electricity demand. Having used exponential smoothing, what advanced method should I consider for more accurate demand predictions?
44. In healthcare, I'm analyzing patient vital signs over time. What advanced time series model would enhance the predictive accuracy of health outcomes?
45. As a data scientist in transportation, I'm forecasting passenger flow. Beyond traditional methods, which advanced technique should I adopt to improve forecasting accuracy?
46. In the context of fraud detection, I'm evaluating the performance of my models. Beyond accuracy and precision, what other evaluation metrics should I consider for a more comprehensive assessment?
47. As a data scientist in healthcare, I'm evaluating models for disease prediction. Which evaluation method would help in understanding the trade-offs between sensitivity and specificity more effectively?
48. Working in marketing, I'm assessing the effectiveness of customer segmentation models. Beyond silhouette score, what other evaluation techniques should I use to validate the quality of my clusters?
49. In the field of education, I'm evaluating models for predicting student success. What advanced evaluation metrics should I consider to ensure my models are both accurate and fair?
50. As a data scientist in manufacturing, I'm assessing the performance of predictive maintenance models. Beyond traditional metrics, which evaluation method should I use to measure the real-world impact of my models?

B. 5 Textbooks

1. Aggarwal, C. C. (2018). *Neural Networks and Deep Learning: A Textbook*. Springer.
2. Alpaydm, E. (2014). *Introduction to Machine Learning* (3rd ed.). The MIT Press.
3. Bruce, P., & Bruce, A. (2017). *Practical Statistics for Data Scientists: 50 Essential Concepts*. O'Reilly Media.
4. Langr, J., & Bok, V. (2019). *GANs in Action: Deep Learning with Generative Adversarial Networks*. Manning Publications.
5. Montgomery, D. C., Jennings, C. L., & Kulahci, M. (2015). *Introduction to Time Series Analysis and Forecasting* (2nd ed.). Wiley.

C. 17 Textbooks

1. Aßenmacher, Matthias. *Multimodal Deep Learning*. Self-published, 2023.
2. Bertsekas, Dimitri P. *A Course in Reinforcement Learning*. Arizona State University.
3. Boykis, Vicki. *What are Embeddings*. Self-published, 2023.
4. Bruce, Peter, and Andrew Bruce. *Practical Statistics for Data Scientists: 50 Essential Concepts*. O'Reilly Media, 2017.
5. Daumé III, Hal. *A Course in Machine Learning*. Self-published.
6. Deisenroth, Marc Peter, A. Aldo Faisal, and Cheng Soon Ong. *Mathematics for Machine Learning*. Cambridge University Press, 2020.
7. Devlin, Hannah, Guo Kunin, Xiang Tian. *Seeing Theory*. Self-published.
8. Gutmann, Michael U. *Pen & Paper: Exercises in Machine Learning*. Self-published.
9. Jung, Alexander. *Machine Learning: The Basics*. Springer, 2022.
10. Langr, Jakub, and Vladimir Bok. *Deep Learning with Generative Adversarial Networks*. Manning Publications, 2019.

11. MacKay, David J.C. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
12. Montgomery, Douglas C., Cheryl L. Jennings, and Murat Kulahci. *Introduction to Time Series Analysis and Forecasting*. 2nd Edition, Wiley, 2015.
13. Nilsson, Nils J. *Introduction to Machine Learning: An Early Draft of a Proposed Textbook*. Stanford University, 1996.
14. Prince, Simon J.D. *Understanding Deep Learning*. Draft Edition, 2024.
15. Shashua, Amnon. *Introduction to Machine Learning*. The Hebrew University of Jerusalem, 2008.
16. Sutton, Richard S., and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd Edition, MIT Press, 2018.
17. Alpaydin, Ethem. *Introduction to Machine Learning*. 3rd Edition, MIT Press, 2014.

D. Sample Output For Baseline Model for Question 40

Question	Contexts	Answer	Context Relevance	Faithfulness	Answer Relevance
As a data scientist in social media analysis, I'm detecting trends from tweets. Beyond LSTM, which cutting-edge NLP technique should I consider for more insightful trend analysis?	<ul style="list-style-type: none"> ample, in the 2016 U.S elections, along with #Election2016, the two political parties used the #VoteTrump and #votedforHillaryClinton respectively to reach out to the general public on social media. In Twitter, for a hashtag to qualify as a trending topic, the Twitter trending algorithm is designed in such a way that the users will be shown the trending tweets based on their location, interests and the profiles they follow. Trends are usually the hashtags that have been used vastly in a short period. Facebook also has a trending module to show what topics are being discussed. Trends are an adequate way to ignite a discussion on a topic. However, considering the dynamic and erratic nature of social media, a trending topic can be replaced within seconds of its initiation; this can have a disadvantage over a current problem, which is trending. If the issue at hand carries a significant concern and needs to reach out to its con- The engagement for Vauxhall remained at a consistent level throughout the event (see Figures 5 and 11), despite the rapid drop-off in the use of the promoted hashtag. CONCLUSIONS & FUTURE DIRECTIONS In this paper, we present a measurement-driven study of the effects of promoted tweets and trends on Twitter on the engagement level of users, using a number of ML and NLP techniques in order to detect relevant tweets and their sentiments. Our results indicate that the use of accurate methods for sentiment analysis, and robust filtering for topical content, is crucial. Given this, we then see that promoted tweets and trends differ considerably in the form of engagement they produce and the overall sentiment associated with them. We found that promoted trends lead to higher engagement volumes than promoted tweets. However, although promoted tweets obtain less engagement than promoted trends, their engagement forms are often more brand-inclusive. features of the event as the location of the event's peak, its growth and relaxation signatures, contain valuable information and must be used in the analytics. The predictive power of Twitter events in this study was evaluated using two scenarios: in the first scenario, we performed the analysis for the aggregated Twitter signal considering only tweets' sentiment; in the second scenario, we cluster Twitter sentiment events based on their shape and then calculate the statistics of successful predictions separately for every event type. The results of our analysis can be summarized as follows: - Aggregated Twitter events can be used to predict sales spikes. - Spikes can be separated into categories based on their shapes (position of the peak, growth, and relaxation signatures). - Different spike shapes are differently associated with sales. - Some spike types have a higher predictive power than the aggregated Twitter signal. This shows that Twitter users do not tweet as positively about a promoted trend as they would about a promoted tweet. Instead, a large proportion of tweets relating to a promoted trend contained no emotional words, or if they did, the positive and negative sentiments balanced each other out. They generally contained just the promoted hashtag or generally had an objective, matter-of-fact tone (e.g., - "Get 3G where I live... #O2WhatWouldYouDo").Figure 11:# related engagements Social media sites like Facebook and Twitter have emerged as popular destinations for users to get real-time news about the world around them. On these sites, users are increasingly relying on crowdsourced recommendations called Trending Topics to find important events and breaking news stories. Typically topics, including key phrases and keywords (e.g., hashtags), are recommended as trending when they exhibit a sharp spike in their popularity, i.e., their usage by the crowds suddenly jumps at a particular time (Twitter 2010). Once a topic is selected as trending, it gets prominently displayed on social media homepages, thus reaching a large user population. Additionally, traditional news organizations often pick trending topics and cover them in their news stories (a practice termed Hashtag journalism), further amplifying their reach. Recognizing their importance, Benevenuto, F., Magno, G., Rodrigues, T., and Almeida, V. 2010. "Detecting spammers on Twitter," in Collaboration, electronic messaging, anti-abuse and spam conference, pp. 12-21. Bermingham, A., and Smeaton, A. F. 2011. "On using Twitter to monitor political sentiment and predict election results," in Proceedings of the Sentiment Analysis where AI meets Psychology Workshop. Bollen, J., Mao, H., and Zeng, X. 2011. "Twitter mood predicts the stock market," Journal of Computational Science (2:1), pp. 1-8. Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, P. K. 2010. "Measuring User Influence in Twitter: The Million Follower Fallacy," in Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media, pp. 10-17. Culotta, A. 2013. "Lightweight methods to estimate influenza rates and alcohol sales volume from Twitter messages," Language resources and evaluation (47:1), pp. 217-238. Dann, S. 2010. "Twitter content classification," First Monday (15:12). 	<p>Consider using transformer-based models like BERT or GPT for more insightful trend analysis in social media. These models excel in understanding context and semantics, which can enhance sentiment analysis and topic detection. Additionally, exploring techniques like clustering for sentiment events can provide deeper insights into the dynamics of trending topics.</p> <p>REFERENCES</p> <ul style="list-style-type: none"> •Poddar, D. (2020, April 23). War of the hashtags: Trending new hashtags to override critical topics in social media. •Dacres, S., Haddadi, H., and Purver, M. (2013, December 23). Topic and sentiment analysis on OSNs: A case study of advertising strategies on Twitter. •Kolchyna, O., Souza, T. T. P., Aste, T., and Treleven, P. C. (2015, August 17). In quest of significance: Identifying types of Twitter sentiment events that predict spikes in sales. •Chakraborty, A., Messias, J., Benevenuto, F., Ghosh, S., Ganguly, N., and Gummadi, K. P. (2017, April 1). Who makes trends? Understanding demographic biases in crowdsourced recommendations. •Dijkman, R., Ipeiro-tis, P., Aertsen, F., and van Helden, R. (2015, March 16). Using Twitter to predict sales: A case study. 	0.0059	0.25	0.8943

E. Sample Output For the Most Enhanced Model For Question 40

Question	Contexts	Answer	Context Relevance	Faithfulness	Answer Relevance
As a data scientist in sports analytics, I'm clustering player performance metrics. Beyond Gaussian Mixture Models, which advanced clustering method should I consider for improved insights?	<p>Global lifestyle and investment patterns are significantly influenced by socioeconomic changes. The widespread availability of high-speed internet and online social media has enabled people to express their views freely on various scales. Nowadays, a majority of users utilize platforms such as WhatsApp, Twitter, Instagram, etc., to communicate and share comments through tweets, status updates, stories, reels, shorts, etc. The digital revolution has also transformed the financial market, making stocks, mutual funds, and precious metals easily accessible in electronic form, allowing people to make online purchases conveniently from the comfort of their homes. Throughout the last century, economics has achieved unprecedented heights. In the multicurrency system, nearly all governments worldwide strive to manage their country's economic status by boosting GDP, reducing inflation, and controlling the exchange rate between their domestic currency and foreign currency to enhance the lifestyle and income levels of their citizens. Microeconomics focuses on the financial position of entities such as farms and enterprises, guiding decision-making and choice among various options. Stocks, mutual funds, oil, gas, precious metals, goods, services, and products witness fluctuations in their market prices over time. These variations are influenced by factors such as news, natural calamities, political stability, climate, and other dynamic elements. The nature of news has changed considerably, with users now primarily accessing information online through various social media platforms rather than relying solely on traditional printed newspapers. Consequently, social media posts play a crucial role in shaping the financial market, especially when authored by influential leaders, significantly amplifying their impact. In this study, we explore the impact of social media posts, such as tweets, on the financial market using artificial intelligence. We analyze the relationship between tweets and stock prices through the application of Natural Language Processing. This motivation propels us to develop a sophisticated model that can effectively analyze the influence of social media posts, particularly tweets, on financial markets.</p> <p>1. Beyond Traditional Models and Temporal Sentiment Dynamics: In contrast to existing financial market prediction models employed by investment banks, which predominantly rely on historical price analysis, our model extends its reach by considering the influential voices of market leaders on social media platforms. By incorporating Twitter sentiments, we introduce a state-of-the-art Memory-based Sentiment Analysis, which goes beyond immediate effects by gauging the temporal impact of tweets on financial market prices. Twitter Sentiment: Our model adopts a multistep learning process, beginning with learning to analyze Twitter sentiment. This phase involves the generation of sentiment scores and categorization into positive, negative, or neutral sentiment statuses based on leaders' tweets. This state-of-the-art approach enables the system to discern the emotional tone of leaders' voices in real-time. Additionally, historical sentiment patterns from previous n days are used to capture more enduring effects on market prices. This dual approach enhances our model's predictive accuracy by considering both the immediate and longer-term impacts of social media sentiments on financial markets. Building on sentiment analysis, our model progresses to learn to map tweet sentiments with different financial components. Here, the model establishes a connection between sentiment scores and statuses with dated market prices, effectively synchronizing the ebb and flow of sentiments with market trends. This mapping provides a comprehensive understanding of how social media sentiments correlate with financial components. Mapping: The final phase, learning to predict the price using tweets stock mapping, equips the model with the ability to forecast financial market fluctuations based on the established mapping between social media sentiments and historical prices. This integration allows our model to offer timely and state-of-the-art predictions that consider both traditional market factors and social media dynamics. To enhance interpretability, this essentially teaches a backtraversal mechanism. By analyzing the effects of specific words or sets of words on Twitter that influence market fluctuations, the system gains the ability to identify and highlight the linguistic triggers influencing financial markets. This granular analysis enhances the model's transparency and provides state-of-the-art insights into the relationship between social media sentiments and market movements.</p>	<p>Consider using Memory-based Sentiment Analysis, which captures both immediate and longer-term impacts of tweets on financial markets. This technique allows for the generation of sentiment scores and the mapping of these sentiments to market prices, enhancing predictive accuracy. Additionally, it provides insights into the linguistic triggers that influence market fluctuations, improving model interpretability.</p> <p>REFERENCES</p> <p>• Das, A., Nandi, T., Saha, P., Das, S., Mukherjee, S., Naskar, S. K., and Saha, D. (2024). Effect of Leaders' Voice on Financial Market: An Empirical Deep Learning Expedition on NASDAQ, NSE, and Beyond.</p>	1.0	0.75	0.8618