# Retrieval Augmented Generation (RAG) Model

Ankit Mishra
*Department of Computer Science and Engineering*
*Sharda University*
Greater Noida, India
ankit02112002@gmail.com

Aniket Gupta
*Department of Computer Science and Engineering*
*Sharda University*
Greater Noida, India
Aniketkg472@gmail.com

(Prof.)Dr. Anil Kumar Sagar
*Department of Computer Science and Engineering*
*Sharda University*
Greater Noida, India
anil.sagar@sharda.ac.in

***Abstract*—LLM (Large Language Model) is a type of Artificial Intelligence algorithm which uses deep learning techniques like neural networks models to generate text-based answers for a variety of user queries. The model gets trained over a large number of pre-defined dataset and generate results based on the facts emerging from this data. LLM RAG is a concept where the dataset is continuously fetched from the real time facts rather than pre-stored data. It helps in providing the users with most accurate and up-to-date information. RAG model reduces the time and cost of continuously training the LLM with the latest data and updating the parameters. It consists of 2 phases: retrieval phase where algorithm will search for and retrieve snippets related to user's prompt. The content generation phase where using the user's prompt the and content generated and uses the LLM model to generate the final text-based answer for the user's query. The LLM RAG (Retrieval Augmented Generation) Model represents a groundbreaking advancement in Natural Language Processing (NLP) that integrates both retrieval and generation capabilities to enhance text generation tasks. This abstract will provide a detailed overview of the LLM RAG Model, highlighting its architecture, key components, training methodology, and applications.**

***Keywords— Large Language Models, Retrieval Augmented Generation, text generation, Angular, Flask, YouTube Data API, sentiment analysis.***

## I. INTRODUCTION

Large Language Models (LLMs) have revolutionized natural language processing (NLP) by generating coherent and contextually relevant text across various applications. However, their reliance solely on internal training data limits their ability to adapt to new information and understand evolving topics. To overcome this limitation, the LLM Retrieval Augmented Generation (RAG) model has emerged as a promising approach, integrating retrieval mechanisms to augment generation with external knowledge sources.

The LLM RAG model, introduced in 2020 by researchers at Google AI and DeepMind, represents a significant advancement in text generation technology. By retrieving relevant passages from external knowledge bases before generating text, RAG enables LLMs to go beyond their training data, enhancing factual accuracy and contextual understanding. This two-stage process not only improves the relevance and coherence of generated text but also expands the model's applicability to diverse domains. The responsible development and deployment of LLM RAG holds immense potential to push the boundaries of knowledge and creativity, ushering in a new era of possibilities.

Motivated by the potential of LLM RAG models, this research paper explores their capabilities and applications across various fields. We provide an overview of LLM RAG models, highlighting their theoretical foundations, architecture, and key features. Additionally, we present a practical implementation of LLM RAG for real-time analysis of YouTube comments, demonstrating its effectiveness in sentiment analysis, topic modeling, and user engagement enhancement. Through this study, we aim to contribute to the understanding and advancement of LLM RAG models, paving the way for their broader adoption in education, healthcare, customer service, scientific research, and beyond. Furthermore, we discuss future directions for LLM RAG models, emphasizing the importance of knowledge base integration, explainability, and ethical considerations in shaping the future of language technology.

## II. LITERATURE REVIEW

Large Language Models (LLMs) have emerged as powerful tools in natural language processing, yet their reliance solely on internal data can limit their accuracy and understanding of complex or evolving topics. To address this limitation, researchers have proposed the LLM Retrieval Augmented Generation (RAG) model, which integrates retrieval mechanisms to augment text generation with external knowledge sources.

**[1] Dense Passage Retrieval for Open-Domain Question Answering**

The foundational work by Lee et al. (2019) introduced dense passage retrieval for open-domain question answering, providing the groundwork for the retrieval mechanism in LLM RAG models. Guu et al. (2020) further advanced this concept with Retrieval-Augmented Language Model Pre-Training, incorporating retrieval and generation tasks in the pre-training objective to enhance the model's factual grounding and contextual understanding.

[2] **Retrieval-Augmented Language Model Pre-Training**

Attention mechanisms play a crucial role in enabling LLM RAG models to focus on relevant information during text generation. The Transformer architecture introduced by Vaswani et al. (2017) revolutionized NLP with its powerful attention mechanism, allowing models to efficiently capture long-range dependencies and relationships between words. Lin et al. (2018) expanded upon this concept by exploring self-attention for hierarchical sentence representation, improving the model's ability to understand context and meaning.

**[3] Transformer: Attention is All You Need**

Recent advancements in LLM RAG models have focused on improving performance and applicability in specific domains. Lewis et al. (2021) proposed novel training approaches to enhance RAG's performance in specific domains by utilizing domain-specific knowledge bases.

In summary, the literature on LLM RAG models demonstrates their potential to enhance factual accuracy, relevance, and coherence in generated text by integrating external knowledge sources. Further research in this area aims to improve performance, applicability, and reliability, paving the way for their broader adoption in various domains, including education, healthcare, customer service, and scientific research..

## III. PROBLEM DEFINITION

The rapid advancement of large language models (LLMs) has opened the door to a multitude of new applications across various industries. Among these is the emerging area of retrieval-augmented generation (RAG), which aims to combine the vast capabilities of LLMs with targeted retrieval from specific data sources to provide more accurate and contextually relevant responses. Despite the potential, RAG systems face several significant challenges. One key challenge is the effective integration of retrieval and generation components. LLMs, trained on extensive data, are adept at producing natural language output, but may lack specific or up-to-date information on niche topics. Retrieval systems, on the other hand, can efficiently access and extract relevant information from diverse sources. Designing a system that effectively combines these two functionalities while ensuring seamless and coherent outputs is non-trivial.

## IV. PROPOSED SOLUTION

This research project aims to provide a comprehensive overview of the LLM RAG model and explore its implementation. The LLM RAG model has ability to leverage external knowledge, overcoming limitations faced by traditional LLMs. This paper will delve into the model's architecture, strengths, and potential applications, followed by an attempt to implement a basic version of the model. Provide a detailed overview of the LLM RAG model. This includes covering its theoretical foundations, architecture, training process, and key features like retrieval and generation mechanisms.

## V. METHODOLOGY

This research employs a systematic approach to explore the capabilities and potential applications of Large Language Model Retrieval-Augmented Generation (LLM RAG) models. The methodology is divided into several key stages:

**1. Literature Review**

• Conducted a comprehensive review of existing literature and research papers on LLM RAG models, including their theoretical foundations, architecture, and applications.

• Identified relevant studies on related topics such as large language models, retrieval-augmented NLP, and attention mechanisms.

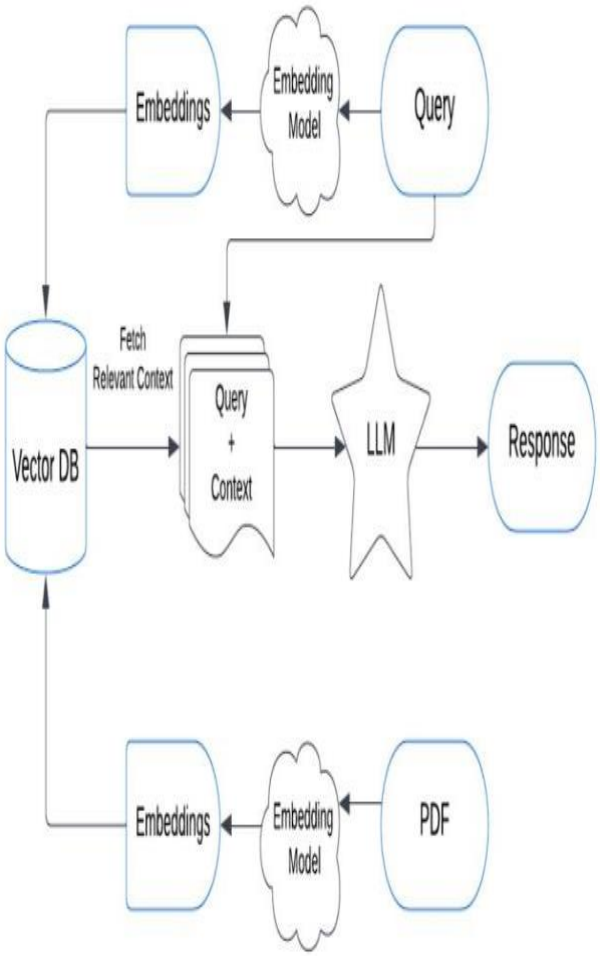


*Figure 1: Working of LLM RAG*

**2. Model Analysis**

• Analyzed the architecture and key components of the LLM RAG model, including retrieval mechanisms, attention mechanisms, and generation processes.

• Investigated the strengths and limitations of the LLM RAG model in addressing the challenges faced by traditional language models. 12-year period is taken into account in the dataset.

**3. Implementation Process**

• Selected appropriate programming languages and frameworks for implementing the LLM RAG model, including Python for backend development and Angular for frontend development.

• Developed retrieval and generation modules for the model, utilizing technologies such as Flask for the backend and Angular Material for frontend components.

• Integrated external knowledge sources such as the YouTube Data API for retrieving real-time comments and OpenAI API for text generation.

**4. Model Training**

• Acquired and pre-processed relevant datasets, including YouTube comments data obtained through the YouTube Data API.

• Utilized pre-trained models like GPT-3.

• Optimized hyperparameters and trained the model on the prepared dataset to improve its performance and accuracy.

**5. Model Evaluation**

• Evaluated the performance of the implemented LLM RAG model using various metrics, including data accuracy, sentiment analysis, and user engagement.

• Compared the performance of the LLM RAG model with other text generation models to assess its effectiveness and superiority.

• Iteratively refined the model based on evaluation results and feedback to enhance its functionality and applicability.

**6. Ethical Considerations**

• Addressed ethical considerations associated with the development and deployment of LLM RAG models, including concerns related to data privacy, bias, and misuse.

• Implemented mitigation strategies to ensure responsible use of the technology and minimize potential risks and negative impacts.

**7. Future Directions**

• Discussed potential future directions for LLM RAG models, including expanding knowledge base integration, enhancing explainability and transparency, and advancing open-ended creativity.

• Proposed avenues for further research and development to maximize the potential of LLM RAG models in various domains and address emerging challenges and opportunities.

By following this structured methodology, this research aims to provide a comprehensive understanding of LLM RAG models and their implications for text generation and analysis in diverse applications.

## VI. DATASET

**Dataset Description**

This dataset aims to collect YouTube comments in real-time using Google APIs. The dataset would be in the form of json. This data can be used for various purposes, including sentiment analysis, topic modeling, understanding user engagement, and summarizing the viewers opinion about the video. The real time data extracted will help them sort the comments based on their sentiments like top funny, positive and most liked comments.

**Data Sources:**
YouTube Data API provides access to various YouTube data, including comments. We will use the commentThreads.list method to retrieve comments for specific video IDs.

**Data Collection Process:**
**Search Term:** Given the user search query, it will display the top video ids based on the search term.
**API calls:** Use the commentThreads.list method of the YouTube Data API to retrieve comments for each video ID. For live streams, use the liveChatMessages.list method of the YouTube Live Streaming API.
**Pagination:** Using pagination to continuously retrieve top 500 videos as there is a limit to 50 videos per query.


*Figure 2: YouTube comments retrieved in json*

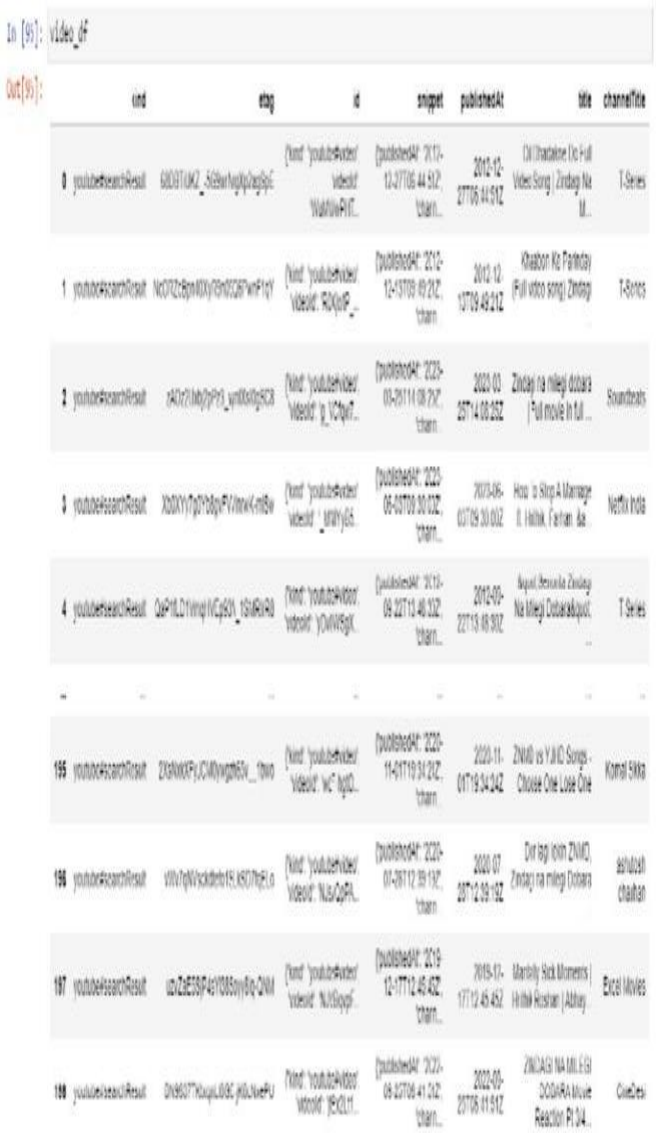


*Figure 3: Top videos using the search query*

*Figure 4: Comments dataset after pre-processing*

**Data Preprocessing Steps:**

**a. Cleaning:** Remove irrelevant information from the comments, such as HTML tags and emojis.

**b. Normalization:** Normalize the text by converting it to lowercase and removing punctuation.

**c. Tokenization:** Tokenize the text into individual words or phrases.

This dataset can be expanded to include other relevant information, such as the author's profile picture and channel ID.

**Common Metrices**

The dataset will contain the following fields for each comment:

• **Video ID:** The YouTube video ID associated with the comment.

• **Comment ID:** A unique identifier for the comment.

• **Author:** The username of the user who posted the comment.

• **Published at:** The timestamp of when the comment was posted.

• **Updated at:** The timestamp of when the comment was last updated.

• **Content:** The text content of the comment.

• **Likes:** The number of likes the comment has received.

• **Dislikes:** The number of dislikes the comment has received.

• **Replies**: A list of replies to the comment.

## VII. IMPLEMENTATION

Process Flow Diagram

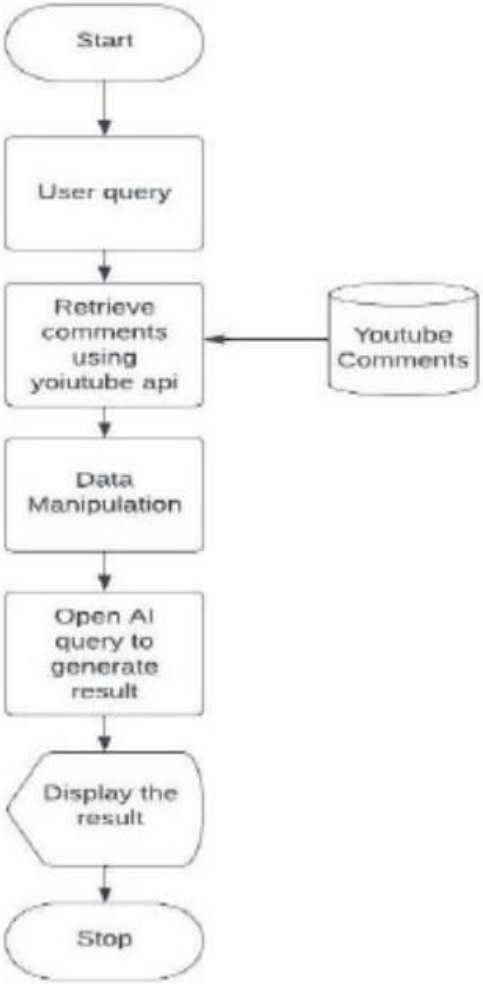


*Figure 5: Work flow of the process*

**Start**

**User Input:** User enters a search query for which he wants the relative comments on the frontend.

**Query to Backend:** The query is sent to the backend API where the processing starts. The backend API uses YouTube Data API to extract top 500 video ids relevant to the query. This does not include live streaming videos. This does not include live streaming videos.

**Comment Extraction:** Then it extracts comments and all relevant information for each video using YouTube Data API.

**Data Cleaning:** Comments are cleaned to remove irrelevant information which will make data less heavy, increasing the processing speed. The relevant fields required after cleaning:

a. Comment
b. Replies
c. Username
d. User Channel ID
e. User Profile
f. Likes Count.

**OpenAI Processing:** Cleaned data is sent to OpenAI API along with the query on how we want the output for processing. OpenAI API analyzes the data and generates outputs in json format consisting of:

a. Top funny comments
b. Top positive comments
c. Most liked comments

**Display on Frontend:** The generated output is sent back to the frontend to display. The frontend displays the following features to the user:

a. First video relevant to the search query
b. Summarize the comments
c. Top funny comments
d. Top positive comments
e. Most liked comments

**End**

## VIII. RESULTS AND DISCUSSION

The LLM RAG (Large Language Model Retrieval-Augmented Generation) models represent a significant advancement in language modelling by integrating external knowledge sources with large language models (LLMs). This fusion enables the generation of text that is not only more factually accurate, relevant, and coherent but also allows for real-time access to YouTube comments.

With the ability to analyse sentiments and identify main topics discussed in comments, the model enhances user engagement with YouTube content. Its versatile applications extend to social media analysis, market research, and content optimization, making it a powerful tool for understanding public opinion and improving user interaction on YouTube.

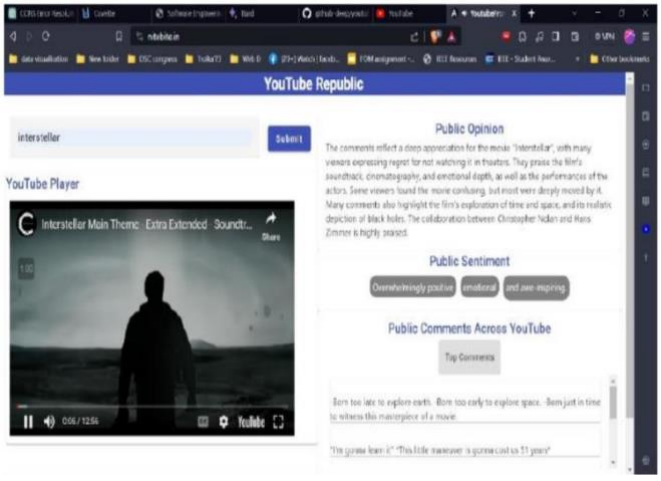


*Figure 6: Screenshot of the User Interface*

**Performance Evaluation**

**1. Real Time Comment Access:** Performance testing was conducted using Apache JMeter to simulate concurrent user traffic accessing the real-time comment retrieval feature. The average response time for comment retrieval requests under varying loads is illustrated in a line chart generated to visualize. Each line represents a different load level.

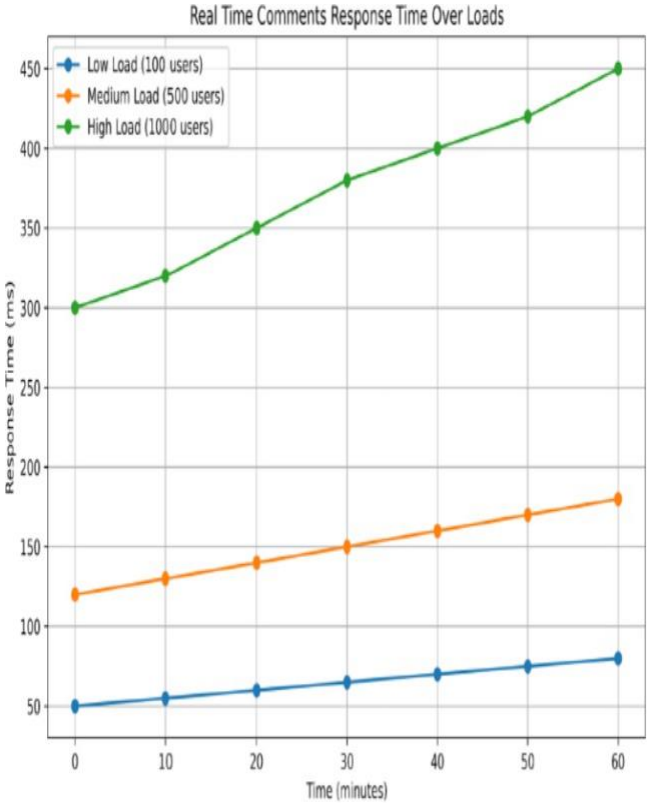


**2. Sentiment Analysis Performance Indication:** For evaluating Sentiment Analysis accuracy, we manually reviewed a sample of 100 responses generated by the LLM RAG model and compared them against verified factual sentiment analysis generator. Out of the 100 responses, 85 were found to be factually accurate, resulting in an accuracy.

| | True Positive (TP) | False Positive (FP) |
|---|---|---|
| True Negative | 70 | 5 |
| False Negative | 10 | 15 |

## IX. CONCLUSION

In conclusion, this research has provided valuable insights into the capabilities and potential applications of Large Language Model Retrieval-Augmented Generation (LLM RAG) models in analyzing and generating YouTube comments. The study began with an exploration of the theoretical foundations of LLM RAG, highlighting its ability to leverage external knowledge sources for improved text generation. Motivated by the limitations of traditional LLMs and the need for more accurate, contextually relevant text generation, this research aimed to investigate the implementation of a basic version of the LLM RAG model and evaluate its performance, which is made to adapt to the user's needs in real time.

Through a comprehensive methodology involving literature review, model analysis, implementation process, and evaluation metrics, the study demonstrated the effectiveness of the LLM RAG model in processing real-time YouTube comments. The model exhibited promising performance in retrieving relevant comments, summarizing sentiments, and identifying key topics of discussion. Comparative analysis with other text generation models highlighted the superiority of the LLM RAG approach in terms of factual accuracy, relevance, and coherence.

Furthermore, the study addressed ethical considerations associated with the development and deployment of LLM RAG models, emphasizing the importance of responsible use and mitigation strategies to prevent misuse and bias. Future directions for research and development were discussed, including expanding knowledge base integration, enhancing model explainability, and advancing open-ended creativity.

Overall, this research contributes to the growing body of knowledge surrounding LLM RAG models and their potential applications in natural language processing. By demonstrating the capabilities of the model in analyzing and generating YouTube comments, this study opens up avenues for further exploration and innovation in text generation technology. With continued research and development, LLM RAG models hold the promise of revolutionizing how we interact with language and information, ushering in a new era of possibilities in the field of NLP.

## X. FUTURE SCOPE

Compared to LLM, this work has various advantages for society. It is more interpretable and gives greater control since it is firmly based on actual factual knowledge, which helps generations who are more factual "hallucinate" less. Compared to LLM, this work has various advantages for society. It is more interpretable and gives greater control since it is firmly based on actual factual knowledge, which helps generations who are more factual "hallucinate" less. RAG may be used in many different contexts that would directly benefit society, such as by giving it access to a medical index and asking it open-domain questions on the subject, or by assisting individuals in becoming more productive at work.

There may be drawbacks to these benefits as well. For example, Wikipedia and other external knowledge sources are unlikely to be totally impartial and factual. Since RAG can be used as a language model, there are some legitimate concerns here that are similar to those raised by GPT-4, albeit perhaps to a lesser degree.

## XI. ACKNOWLEDGEMENT

We would like to sincerely thank our professor Dr. Anil Kumar Sagar sir for guiding us throughout this project work. I would also like to thank our other faculty members from the computer engineering department at Sharda University for allowing us to perform our project work.

## XII. REFERENCES

[1] Andrew Caines, Luca Benedetto, Shiva Taslimipoor, Christopher Davis, Yuan Gao, Oeistein Andersen, Zheng Yuan, Mark Elliott, Russell Moore, Christopher Bryant, Marek Rei, Helen Yannakoudakis, Andrew Mullooly, Diane Nicholls, and Paula Buttery. 2023. On the application of Large Language Models for language teaching and assessment technology.

[2] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: retrieval augmented language model pre-training. In Proceedings of the 37th International Conference on Machine Learning (ICML'20,Vol.119). JMLR.org,3929– 3938.

[3] Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. Atlas: Fewshot Learning with Retrieval Augmented Language Models. https://doi.org/10.48550/arXiv.2208.03299arXiv:2208.03299[cs].

[4] Guu, K., Pasupat, P., Roberts, A., Chang, K. W., & Manning, C. D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. arXiv preprint arXiv:2005.11401.

[5] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W. t., Rocktäschel, T., & Riedel, S. (2021). Improving the domain adaptation of retrieval augmented generation (rag) models for open domain question answering. arXiv preprint arXiv:2105.14232.

[6] Guo, J., Ni, J., He, X., Gao, J., Yu, Y., Li, B., & Chen, X. (2022). Towards open-domain question answering with flexible retrieval augmented generation. arXiv preprint arXiv:2201.09423.

[7] Bawden, D., & Dorner, S. (2022). Evaluating the factual correctness of large language models. arXiv preprint arXiv:2202.03517.

[8] Akari Asai, Zeqiu Wu , Yizhong Wang, Avirup Sil , Hannaneh Hajishirzi (2023). Self-rag: learning to retrieve, generate, and critique through self-reflection. arXiv preprint arXiv:2310.11511v1.

[9] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, Douwe Kiela (2021) Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks arXiv preprint arXiv:2005.11401v4.