

RAG-FUSION: A NEW TAKE ON RETRIEVAL-AUGMENTED GENERATION

Zackary Rackauckas

Power and Sensor Systems, Infineon Technologies, San Jose, California

ABSTRACT

Infineon has identified a need for engineers, account managers, and customers to rapidly obtain product information. This problem is traditionally addressed with retrieval-augmented generation (RAG) chatbots, but in this study, I evaluated the use of the newly popularized RAG-Fusion method. RAG-Fusion combines RAG and reciprocal rank fusion (RRF) by generating multiple queries, reranking them with reciprocal scores and fusing the documents and scores. Through manually evaluating answers on accuracy, relevance, and comprehensiveness, I found that RAG-Fusion was able to provide accurate and comprehensive answers due to the generated queries contextualizing the original query from various perspectives. However, some answers strayed off topic when the generated queries' relevance to the original query is insufficient. This research marks significant progress in artificial intelligence (AI) and natural language processing (NLP) applications and demonstrates transformations in a global and multi-industry context.

KEYWORDS

Chatbot, Retrieval-augmented Generation, Reciprocal Rank Fusion, Natural Language Processing

1. INTRODUCTION

Infineon account managers and field application engineers have expressed the need to obtain sales-oriented product information rapidly, but Infineon's product selection guides and datasheets are often hundreds of pages long. Born from this need was the rapid development of chatbots to provide account managers and engineers with technical information in seconds. These chatbots are built off the state-of-the-art retrieval-augmented generation framework.

Recently, retrieval-augmented generation (RAG) has been at the core of all of Infineon's virtual assistants. Retrieval-augmented generation answers a user's questions related to the purpose of the virtual assistant. The method combines text generation from large language models (LLMs) and document retrieval from databases of related documents to generate accurate, relevant, and comprehensive responses [1]. Large language models are advanced natural language processing systems trained on large sets of data that process and generate text. While they are designed to handle tasks like machine translation, summarization, and conversational interactions, RAG models can also leverage them for information retrieval [2]. RAG has shown remarkable success in several knowledge-intensive natural language processing tasks including accurate trivia question answering, highly accurate fact verification, and accurately answering middle school math questions [3] [4]. In addition, retrieval-augmented generation reduces misinformation typically produced by large language models and non-RAG chatbots [5].

Document retrieval is a fundamental component of the algorithm. Traditional RAG virtual assistants rank documents in the order of relevance to the query, usually by vector distances. This

means that the more relevant a document is in a query, the higher priority it takes being in the answer. Recently, however, developers and researchers have explored implementing different reranking methods for documents. It has been found that reranking in retrieval-augmented generation plays a significant role in improving retrieval results and in increasing the accuracy, relevance, and comprehensiveness of answers [6].

Reciprocal rank fusion (RRF) is a reranking algorithm that gives a reciprocal rank to documents in multiple sources, then combines those ranks and documents into one final reranked list. It has been found that RRF outperforms many other document reranking methods while being sensitive to its parameters [7] [8]. Utilizing RRF as a reranker in a RAG algorithm yields RAG-Fusion, a novel RAG-based chatbot model popularized by Adrian H. Raudaschl [9].

The chatbot I developed in this paper was initially a traditional RAG bot developed to be used by automotive field engineers. But I found that there was a need for customers and distributors to use this chatbot as well from questions asked online in our developer community [10]. Accordingly, I changed the model of the chatbot from RAG to RAG-Fusion with the hypothesis that it would provide answers that not only held the same accuracy as the RAG chatbot but were also more comprehensive in addressing various perspectives of users' questions. Thus, I tested the bot for viability in the following three areas: answering product-specific questions from engineers, answering sales teams' queries about sales strategies, and answering customer queries about products. Specifically, I focused on answering questions related to our silicon MEMs microphones and metal-oxide-semiconductor field-effect transistors (MOSFETs) [11] [12].

2. RAG Vs RAG-FUSION

The traditional retrieval-augmented generation chatbot model for specific product information consists of the following steps:

- Gather product documents (e.g., datasheets, and product selection guides) in a database of documents for retrieval into text.
- Create vector embeddings -- numerical representations of the text that the algorithm can understand -- and store them in a vector database.

Upon receiving a query from the user,

- Retrieve the n most relevant documents based on vector distance to the original query via vector search.
- Send the query together with the retrieved documents to a large language model to generate a response and output the response.

RAG-Fusion, on the other hand, has a few extra steps [13]. Once the original query is received, the model sends the original query to the large language model to generate several new search queries based on the original query. For example, if the user's original query is "Tell me about MEMs microphones," the generated queries may include

- *What are MEMs microphones and how do they work?*
- *What are the advantages of using MEMs microphones?*
- *What are some recommended MEMs microphones?*

The algorithm then performs vector search to find several relevant documents like with RAG. But, instead of sending those documents with the queries to the large language model to generate

the output, the model performs reciprocal rank fusion. Reciprocal rank fusion is an algorithm commonly used in search to assign scores to every document and rerank them according to the scores. The scores assigned to each document, or rrf scores, are

$$rrfscore = \frac{1}{rank + k}$$

where rank is the current rank of the documents sorted by distance, and k is a constant smoothing factor that determines the weight given to the existing ranks.

Upon each calculation of the score, the rrf score is accumulated with previous scores for the same document, and when all scores are accumulated, documents are fused together and reranked according to their scores. The model then sends the reranked results along with the generated queries and the original queries to the large language model to produce an output.

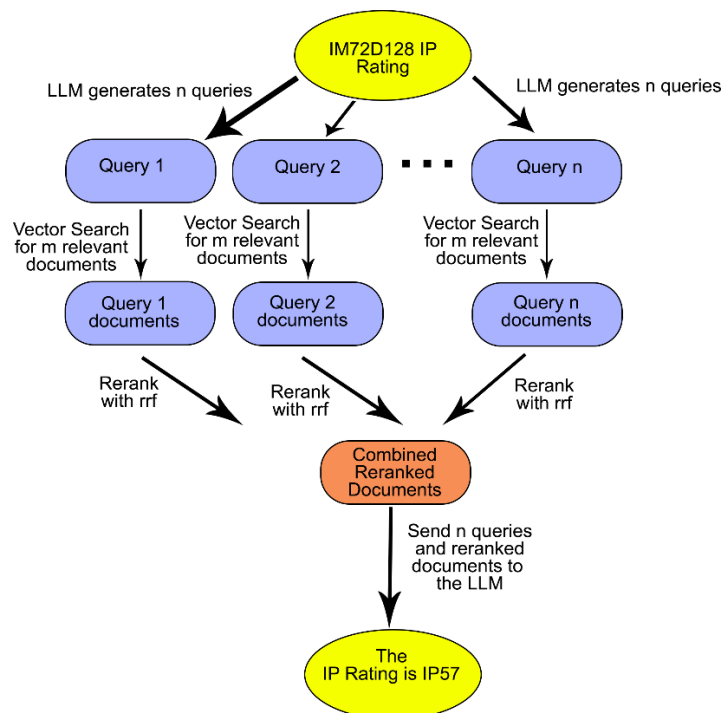


Figure 1. Diagram illustrating the high-level process of RAG-Fusion starting with the original query "IM72D128 IP Rating."

3. INFINEON CHATBOT FOR ENGINEERS

There are three potential use cases of the Infineon RAG-Fusion chatbot: providing technical information to engineers, providing sales-oriented information to account managers, and providing customer-facing information to customers.

To test the viability of the chatbot in providing information to engineers, I gathered a list of technical questions from Infineon's developer community forum regarding our products [10]. I then optimized the questions to produce the best output with the RAG-Fusion. This included correcting spelling and grammar, converting statements into questions, and dividing multi-part questions into a single query. I then compared the accuracy of the RAG-Fusion chatbot in

answering the technical questions against the correct forum solutions provided by Infineon experts. Take the following question for example.

Query: IP rating of mounted IM72D128

Generated queries: ['1. What is the IP rating of the mounted IM72D128?', '2. IP rating explained for mounted IM72D128.', '3. Waterproofing capabilities of the IM72D128 with its IP rating.', '4. How does the IP rating of the IM72D128 affect its durability when mounted?']

Bot answer: The mounted IM72D128 has an IP57 rating, indicating a high level of protection against both dust and water. The first digit, 5, signifies its resistance to solid objects like dust or dirt. The second digit, 7, represents its resistance to water. This rating means that the IM72D128 has a sealed design, preventing water or dust from entering between the membrane and backplate. This robust design ensures no mechanical blockages or electric leakage issues typically seen in MEMS microphones. As a result, the IM72D128 is highly durable and can withstand exposure to water and dust without significant damage. Its sealed MEMS design allows it to function properly even in challenging environmental conditions, providing long-term durability and reliable performance.

Forum solution: Yes, you are right. IM72D128 has IP57 rating. The recommended Reflow soldering and board assembly of the microphone is given in section 7 (page: 14-15) of the datasheet. Please also refer to the application note: MEMS microphone mechanical & acoustical implementation.

In this case, the chatbot provided a correct answer, that the IM72D128 microphone has an IP57 rating. Unlike the forum solution, however, the chatbot not only provided the rating, but also explained what the rating signifies. In addition, it provided information about the microphone's sealed design, ensuring no mechanical blockage or electric leakage issues, and a final statement summarizing its strengths. This extra information was derived from the generated queries and provided an answer that was not only correct, but also more comprehensive and had more context surrounding the question compared to the forum solution. The only part missing from the answer is that the microphone delivers high ingress IP57 protection at a microphone level because of Infineon's latest Sealed Dual Membrane MEMS technology. This missing part may be solved, however, by increasing the number of queries generated or documents retrieved.

In addition to technical product questions, many engineers also ask for help troubleshooting or product instructions. The bot, however, tends to answer these questions with very general instructions and by indicating the user should follow up with Infineon. The following query highlights this tendency.

How to connect and communicate with the PDM interface of IM69D130 with stm32

The bot lists six steps, five of which are mainly general guidance, with one step matching the expert solution:

Bot: *Configure the PDM interface settings on the STM32 microcontroller according to the requirements of the IM69D130, including clock frequency, data format, and other parameters.*

Form: *Set up the data format on the I2S interface to match the PDM data from the IM69D130.*

However, the forum answer emphasizes that the I2S interface of the microcontroller must be in slave receive mode, which is nowhere to be found in the bot's answer. This is most likely because the product database does not contain this kind of guidance. None of the documents mention slave receiver mode, which is a keyword in the response. Since it is unable to retrieve this information, the language model makes up for this by outputting its prior knowledge to supplement the missing information from the documents, which are general troubleshooting steps and not product-specific steps. This tendency can be solved by providing correct documentation in the database regarding troubleshooting and methods to use the product.

4. INFINEON CHATBOT FOR ACCOUNT MANAGERS

Although its original purpose was to be used for engineers, the bot ended up performing the best in response to sales-oriented questions. These are questions that are asked by account managers or inside sales representatives that give insight into the benefits of Infineon products and sales strategies.

To devise sales strategies, the bot retrieves the documents related to the specific products, in the following case our 100V Linear Field-effect Transistor (FET). It then synthesizes the product knowledge with logic from the large language model to create a valuable strategy based on accurate product information. Generating multiple queries is key to devising a comprehensive and multiperspective sales strategy. For example, in response to the prompt

How do I sell a 100V OptiMOS Linear FET for power over ethernet to my customer?

the bot generates unique prompts addressing market trends and customer preferences, best practices for selling electrical components, the value proposition of our Linear FET, and strategies for effectively pitching the benefits and features of our Linear FET. The answer generated by the chatbot addresses these queries by providing technical specifications such as the latest IEEE 802.3bt standard increasing the available power via Power over Ethernet (PoE). It continues to highlight Infineon's expertise in switched-mode power supply designs and our FETs' high efficiency, protection, best-in-class performance with ultra-low RDS(on) and low charge, and high reliability. It responds that the value proposition includes improving efficiency and power density, low conduction losses, and high switching performance leading to energy savings and better thermal management. For sales strategies, it suggests highlighting its value proposition, advantages in performance and efficiency, and the importance to understand customer-specific needs and requirements for their applications. It ends with suggesting the account manager provide technical support and assistance, which ties back into our Infineon chatbot use case for engineers.

This sales-oriented question answering for account managers reduces the need for an account manager to ask an engineer for product information as well as the engineer's need to sift through a 100-page datasheet. This saves times for everyone in the process and increases the account manager's ability to sell the product, resulting in more sales and higher revenue.

Inside sales representatives often need to quickly find product-specific information. This is similar to the use case for engineers, but instead of technical information like voltage and current density, it is information about applications and product validation.

Query: What is the application of IM73A135V01?

Bot answer: Based on the provided documents, the IM73A135V01 is a dust and water resistant analog XENSIV™ MEMS microphone with a best-in-class Signal to Noise Ratio (SNR) of 73dB(A) for high-quality audio pick-up. It also has a flat frequency response and low distortion. In addition, it is IP57 robust to dust and water. While the documents do not specifically mention application examples or industry applications for the IM73A135V01, it is a robust and high-performance MEMS microphone that can be used in various audio devices such as TWS earbuds and compact audio devices. Based on its features and benefits, it can be used in applications that require high-quality audio pick-up, low distortion, and dust/water resistance. Some potential applications could include consumer electronics, automotive, industrial, and smart devices.

5. INFINEON CHATBOT FOR CUSTOMERS

While the chatbot provides great support for sales teams, it is also able to provide comprehensive and helpful feedback for customers' product-related questions. These are not engineering questions, but rather questions from customers who are considering buying our products. For example, a customer could be looking to manufacture a camera and wants to know whether our microphones are a good fit for their camera. In fact, this was a real question asked on our developer community [10].

Query: Is the IM72D128 microphone good to use in an outdoor surveillance camera?

Bot answer: The IM72D128V01 microphone is an ideal choice for outdoor surveillance cameras based on several documents. Document 1 states that it is IP57 dust and water resistant, making it able to withstand the conditions typically encountered in outdoor settings. Document 7 further describes it as an ultra-high-performance digital microphone with a high signal-to-noise ratio (SNR) and low distortions, which ensures clear and high-quality audio capture. This makes it well-suited for outdoor surveillance cameras that require excellent audio capture capabilities. Additionally, Document 8 mentions that the IM72D128V01 microphone is designed for both indoor and outdoor automotive applications, offers best-in-class audio performance, and has an extended temperature range. These features can be advantageous for outdoor surveillance cameras as well, emphasizing its suitability for capturing audio in outdoor environments.

With this answer, the customer understands that our IM72D128 microphone does align with their project, resulting in a higher chance of a new sale.

6. BENEFITS OF A CHATBOT

Customers and engineers tend to ask several questions in one query, and Infineon experts tend to respond to the main question and not the several smaller ones asked. With a chatbot, albeit in separate queries, users can ask all their questions and get answers.

In addition to product-specific knowledge, a chatbot can leverage pretrained knowledge from its large language model that an Infineon expert may not have. For example, one question asked by a customer was regarding whether MEMS microphones could be used in wooden instruments or professional recording studios. While our expert responded by pointing to the MEMS microphone product page and datasheets, the RAG-Fusion bot was able to provide a useful and comprehensive answer drawing on its prior trained knowledge and database of MEMS microphones.

MEMS (Micro-Electro-Mechanical Systems) microphones are small devices that convert sound waves into electrical signals. They utilize a diaphragm that vibrates in response to sound waves and an acoustic sensor that detects these vibrations. The sensor then translates the vibrations into electrical signals, which can be amplified and processed for various applications.

The advantages of using MEMS microphones in stage or professional studio recordings are significant. Their small size allows for easy integration into equipment and placement in different areas. They offer high-performance audio capture with low self-noise, wide dynamic range, and high sensitivity. MEMS microphones also provide excellent audio quality, enabling accurate and precise sound reproduction. Additionally, their compact design enables the creation of microphone arrays for noise cancellation and beamforming, enhancing audio capture in complex acoustic environments.

However, when using MEMS microphones in wooden instruments on stage or in professional studio recordings, there may be limitations or challenges. One potential limitation is their sensitivity to handling noise, as MEMS microphones can be more sensitive to vibrations and movements compared to traditional microphones. This requires careful handling and mounting to minimize unwanted noise. Additionally, the frequency response of MEMS microphones may not perfectly match the tonal characteristics of wooden instruments, necessitating careful selection and placement to achieve the desired sound.

To capture the sound of wooden instruments effectively, some recommended MEMS microphones include models like the Infineon-IM73D122 or the Infineon-IM72D128. These microphones offer high sensitivity, low self-noise, and a wide dynamic range, making them suitable for capturing the nuances and details of wooden instruments. However, it is important to experiment and select the microphone that best suits the specific instrument and desired sound.

In summary, MEMS microphones can be used in wooden instruments on stage or in professional studio recordings, offering advantages such as small size, high audio quality, and the ability to capture details and nuances. However, careful consideration should be given to microphone selection, placement, and handling to overcome any potential limitations or challenges.

7. CHALLENGES OF RAG-FUSION

The largest challenge of RAG-Fusion is the slow answer time from receiving the query to outputting the answer. I compared the runtime to our traditional RAG chatbot by performing ten back-to-back runs with the same query. I then subtracted the time when the query was received from the time the output was given to determine the time it took for that run. Back-to-back runs should control for application programming interfaces (APIs) having different response times at different times of the day.

Table 1. Comparison of RAG-Fusion Time and RAG Time in Smart Speaker Runs. RAG-Fusion takes 1.77 times longer to output an answer.

Run	RAG-Fusion Time (s)	RAG Time (s)
1	42.72	30.48
2	32.05	32.93
3	12.85	25.94
4	42.78	16.70
5	36.58	11.89
6	45.99	10.62
7	34.92	17.58
8	35.56	14.42
9	37.55	28.21
10	25.19	5.44
Average	34.62	19.52

Although it is hard to generalize exact runtime numbers, this average over ten back-to-back runs shows that RAG-Fusion bot is almost certainly slower than the RAG bot. In this case, RAG's average runtime from query to output was 19.52 seconds, with RAG-Fusion being nearly 1.77 times as slow with an average query to output time of 34.62 seconds. Moreover, this trend is evident in all individual runs but runs two and three.

After performing several runs with different queries and at different times, I determined that the slowness of RAG-Fusion compared to RAG can be mostly attributed to the second API call to the large language model. Even in long queries of 70 or more words, the call to generate multiple queries never took more than 5 seconds. The model then ran through document retrieval and reciprocal rank fusion almost instantly until it stopped for the second API call for several seconds. The second API call's complexity is amplified by the volume and diversity of the input in the form of multiple queries and a substantial number of documents in contrast to the much simpler call in RAG with one query and fewer documents. Two ways to overcome this slowness included hosting an LLM locally to reduce latency in the calls to the LLM and decreasing the number of queries generated by the first API call.

Another challenge of RAG-Fusion is the inability to empirically evaluate answers. While evaluation frameworks like ROUGE, BLEU, BLEURT, and METEOR are often applied to assess the accuracy of retrieval-augmented generation models, in the case of the Infineon chatbot, there is not necessarily a direct, expected answer. Thus, such evaluation frameworks are less effective for tasks such as the sales and customer-oriented answering where outputs can vary significantly in structure and content while still being correct.

New evaluation toolkits such as RAGelo and Ragas have recently emerged. Such frameworks seek to provide automated evaluation for RAG-based models. RAGelo takes in a list of questions, documents, and answers. It then assigns a tournament-style Elo ranking of multiple

RAG pipelines over different queries and prompts [14]. Ragas provides empirical assessment scores in context precision, faithfulness, and answer relevancy [15]. Despite promising results on evaluating RAG-based agents, further tweaking to these evaluation methods is needed to align with the unique goals of the Infineon RAG-Fusion chatbot.

The evaluation method that reflects RAG-Fusion's goals are human evaluations based on accuracy, relevance, and comprehensiveness. Ratings from Infineon account managers and engineers provide the best insight into answer quality from an expert lens. For the purposes of this paper, however, I manually evaluated the responses based on the same scale. I found that RAG-Fusion excels in accuracy and comprehensiveness, but occasionally strays in relevance. This occurs when irrelevant queries are generated by the large language model. While the answers to the generated queries were accurate and comprehensive, they strayed from the intention of the original question if the generated query it answered differed too much from the intent of the original query.

To achieve answers that completely match the intention of the original query, users must sometimes engineer prompts to be more specific and comprehensible so that the large language model generates appropriate queries. An example question from a distributor was regarding whether our IM72D128 microphone is good for them to use in a security camera: "If it was Okay, it is very good for out-door surveillance camera," with "If it was Okay" referring to whether water and dust come through the sound port. The first prompt I tried for this question was "Is IM72D128 very good for an out-door surveillance camera?" However, the language model did not recognize my intention of the phrase "good for" and thus generated prompts labelling IM72D128 as a surveillance camera itself. I was able to generate relevant queries with IM72D128 understood to be a microphone used in a surveillance camera with the prompt "Is the IM72D128 microphone good to use in an outdoor surveillance camera?" Situations like this can be remedied by providing a prompt creation guide to users, but in the long run I would like to investigate ways to integrate prompt engineering into the bot side rather than forcing the burden on users.

Another challenge of RAG-Fusion, which is also a challenge of RAG, is its inability to provide definitively negative answers. Below is an example of when the bot could not find documents relating to the user's query.

Query: We need to Confirm whether the IFX waterproof MIC has a Sleeping mode and wake-up functions.

Bot answer: Based on the provided documents, it is uncertain whether the IFX waterproof MIC has sleeping mode and wake-up functions, as none of the documents explicitly mention these features. To obtain accurate information, it is recommended to refer to the official documentation or reach out to the manufacturer directly for confirmation.

Forum answer: Our microphone does not come with sleep mode and wake up functionality. With respect power optimization, our microphone comes with low power mode.

While the bot was correct that there is no mention of sleeping mode and wake-up functions in the database of product information, it did not provide a firmly negative answer. This is because the nature of the bot is centred around document retrieval and not providing clear negative answers.

8. CONCLUSION

I found that the Infineon RAG-Fusion chatbot was able to provide more accurate and comprehensive answers than traditional RAG models. This includes answers for questions related to technical product information for engineers, sales strategies for account managers, and customer-oriented product explanations. This is due to the generation of multiple queries based on the original query, which allows the chatbot to answer questions from multiple angles. In addition, using RRF to rerank the documents retrieved results in the most relevant documents getting the highest priority in answer generation.

While RAG-Fusion increases answer quality, it comes with challenges such as a longer runtime than other models due to a more complex call to the LLM with multiple queries and more documents, answers going off track from the original query due to irrelevant queries generated by the first LLM call, and the occasional need for appropriate prompt engineering to generate the desired outcome.

Many questions asked on the developer forum were from customers and distributors in non-English speaking countries. When customers translate their original questions to English, some context may be lost and thus, customers may not receive the answers they want from Infineon experts. This was clear with the previously discussed question regarding the usage of MEMs microphones in security cameras. In the future, I will expand the Infineon RAG-Fusion chatbot to answer in languages other than English, especially Japanese and Mandarin Chinese.

I will also research ways to better represent multimodal pdf datasheets as text for retrieval-augmented generation. This will allow RAG-Fusion chatbots to retrieve answers from more granular text, resulting in a higher accuracy of specific product information and reduced hallucination frequency. I will also evaluate and tweak systematic frameworks like RAGelo and Ragas according to RAG-Fusion's accuracy and performance. Other future research includes optimizing methods to improve real-time performance optimization, automated quality assurance, and integration into internal and external web platforms.

ACKNOWLEDGEMENTS

The author would like to thank Brooks Felton, Cynthia Meah, and Christopher Arnold.

REFERENCES

- [1] w. yu, "retrieval-augmented generation across heterogeneous knowledge," (2022) proceedings of the 2022 conference of the north american chapter of the association for computational linguistics: human language technologies: student research workshop, pp. 52-58, doi: 10.18653/v1/2022.naacl-srw.7
- [2] h. naveed, u. a. khan, s. qiu, m. saqib, s. anwar, m. usman, et al., (2023) "a comprehensive overview of large language models," doi: 10.48550/arxiv.2307.06435.
- [3] z. levonian, c. li, w. zhu, a. gade, o. henkel, m. postle, et al., (2023) "retrieval-augmented generation to improve math question-answering: trade-offs between groundedness and human preference," neurips'23 workshop on generative ai for education, doi: 10.48550/arxiv.2310.03184.
- [4] p. lewis, e. perez, a. piktus, f. petroni, v. karpukhin, n. goyal, et al., (2020) "retrieval-augmented generation for knowledge-intensive nlp tasks," neurips, doi: 10.48550/arxiv.2005.11401.
- [5] k. shuster, s. poff, m. chen, d. kiela, j. weston, (2021) "retrieval augmentation reduces hallucination in conversation," doi: 10.48550/arxiv.2104.07567
- [6] m. glass, g. rossiello, m. f. m. chowdhury, a. r. naik, p. cai, and a. gliozzo, (2022) "re2g: retrieve, rerank, generate," naacl, doi: 10.48550/arxiv.2207.06300

- [7] g. v. cormack, c. l. a. clarke, and s. buettcher, (2009) "reciprocal rank fusion outperforms condorcet and individual rank learning methods," sigir '09: proceedings of the 32nd international acm sigir conference on research and development in information retrieval, pp. 758-759, doi: 10.1145/1571941.1572114.
- [8] s. bruch, s. gai, and a. ingber, (2023) "an analysis of fusion functions for hybrid retrieval," acm transactions on information systems, vol. 42, no. 20, pp. 1-35, doi: 10.1145/3596512.
- [9] a. h. raudaschl, (2023) "forget rag, the future is rag-fusion,"towards data science.
- [10] "infineon developer community," [online]. available: <https://community.infineon.com/>.
- [11] "xensiv™ – sensing the worldsensor solutions for automotive, industrial, consumer and iot applications," [online].available: <https://www.infineon.com/cms/en/product/sensor/mems-microphones/>.
- [12] "power mosfet - infineon technologies," [online]. available: <https://www.infineon.com/cms/en/product/power/mosfet/>.
- [13] a. h. raudaschl, github - rag-fusion: the next frontier of search technology. [online]. available: <https://github.com/raudaschl/rag-fusion>.
- [14] camara, and f. r. barrera, github – ragelo. [online] available: <https://github.com/zetaalphavector/ragelo>.
- [15] s. es, j. james, l. espinosa-anke, and s. shockaert, (2023) “ragas: automated evaluation of retrieval augmented generation,” doi: 10.48550/arxiv.2309.15217

AUTHOR

Zackary Rackauckas is an AI and businessdevelopment research intern at Infineon Technologies, San Jose, United States. His current interests include multilingual NLP, generative AI, and digitization and decarbonization methods. He is a professional member of IEEE Intelligent Systems. Zack received his BA in mathematics from Swarthmore College. Contact him at zackary.rackauckas@infineon.com.

