

# ASRank: Zero-Shot Re-Ranking with Answer Scent for Document Retrieval

Abdelrahman Abdallah, Jamshid Mozafari, Bhawna Piryani, Adam Jatowt

University of Innsbruck

{abdelrahman.abdallah, jamshid.mozafari, bhawna.piryani, adam.jatowt}@uibk.ac.at

## Abstract

Retrieval-Augmented Generation (RAG) models have drawn considerable attention in modern open-domain question answering. The effectiveness of RAG depends on the quality of the top retrieved documents. However, conventional retrieval methods sometimes fail to rank the most relevant documents at the top. In this paper, we introduce ASRANK<sup>1</sup>, a new re-ranking method based on scoring retrieved documents using zero-shot answer scent which relies on a pre-trained large language model to compute the likelihood of the document-derived answers aligning with the answer scent. Our approach demonstrates marked improvements across several datasets, including NQ, TriviaQA, WebQA, ArchivalQA, HotpotQA, and Entity Questions. Notably, ASRANK increases Top-1 retrieval accuracy on NQ from 19.2% to 46.5% for MSS and 22.1% to 47.3% for BM25. It also shows strong retrieval performance on several datasets compared to state-of-the-art methods (47.3 Top-1 by ASRANK vs 35.4 by UPR by BM25).

## 1 Introduction

Document retrieval is a core sub-task in many NLP problems, including open-domain question answering (ODQA), where a document is retrieved and then read to answer an input query. This process tries to find the most relevant documents or passages given the query. The Retrieval-Augmented Generation (RAG) model has achieved a significant improvement in the field of ODQA (Lewis et al., 2020). RAG models combine retrieved documents and advanced pre-trained large language models (LLMs) generating responses based on the retrieved information (Lewis et al., 2020; Lála et al., 2023). However, the performance of RAG models depends on the top retrieved documents, especially on the first document (Setty et al., 2024;

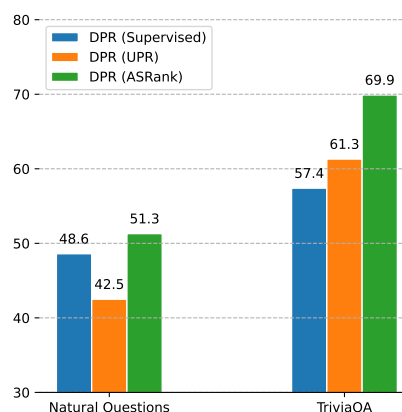


Figure 1: After re-ranking the top 1,000 passages retrieved by DPR (Karpukhin et al., 2020) with ASRANK, our method surpasses the performance of strong unsupervised models like UPR (Sachan et al., 2022) on the Natural Questions and TriviaQA datasets.

Zhang et al., 2024). The RAG model usually uses the first retrieved document, which is the primary source for generating the response. In RAG, queries and documents are embedded in a shared representation space to enable efficient search before using a task-specific model to perform a deeper, token-level document analysis.

In this paper, we introduce ASRANK, a simple, effective, fast, and cost-efficient re-ranking method that leverages the concept of *answer scent* which is analogous to the way in which animals track the scent of their food or prey (Maxwell and Azzopardi, 2018). Cognitive psychologists (Winerman, 2012) have found that people search for information online in much the same way as animals hunt for food, leading to the establishment of the concept of *information scent* in the Information Retrieval field. It refers to the trail of relevant information that leads a user to the correct answer. Our proposal is built upon a similar concept of tracing the

<sup>1</sup><https://github.com/DataScienceUIBK/rankify>

answer scent. It first utilizes larger LLMs like GPT-3.5 or Llama 3-70B to generate an answer scent. This is done just once, hence is computationally efficient. Subsequently, a smaller model such as T5 is employed to re-rank the documents based on the received answer scent. This two-tiered approach allows leveraging the generative capabilities of a larger LLM to boost the re-ranking capabilities of smaller models thanks to improved contextual understanding. Our method scores retrieved documents using a zero-shot answer scent, which relies on a pre-trained LLM to compute the likelihood of the document-derived answers aligning with the answer scent. This approach allows to rank documents not just based on their initial retrieval scores but also on the likelihood of containing an answer (via answer scent) and the degree to which they contain information that aligns with the expected answer. By applying a cross-attention mechanism to every token in both the question and the passage, ASRANK tracks the answer scent within the document corpus. Our approach successfully addresses the challenge of ensuring that the most relevant document is ranked at the top, which is a key component in open-domain question answering and RAG systems (Figure 1).

## 2 Method

In this section, we detail the methodology of ASRANK, starting with retrieving documents based on either sparse or dense techniques. Subsequently, we introduce our concept of generating an Answer Scent using a large language model (Section 2.2), followed by an efficient re-ranking process that employs a smaller model (Section 2.3), which enhances the alignment and relevance of the retrieved documents to the query in our RAG system. Figure 2 shows an overview of the ASRANK framework.

### 2.1 Retriever

Let  $\mathcal{M} = \{d_1, \dots, d_M\}$  represent a collection of evidence documents. Given a query  $q$ , the retriever’s task is to select a subset of relevant documents  $\mathcal{D} \subset \mathcal{M}$ , aiming to include those that likely contain the answer to  $q$ . Our framework is designed to operate on documents retrieved by arbitrary methods, hence ones that can either utilize sparse or dense representations. **Sparse representation** methods such as BM25 (Robertson et al., 2009), a non-neural approach, typically rely on term frequency and inverse document frequency to

rank documents. This method is effective for scenarios where lexical matching is crucial, providing a strong baseline due to its simplicity and proven efficiency in various information retrieval tasks. **Dense representation** methods like Dense Passage Retrieval (DPR) (Karpukhin et al., 2020) employs neural network architectures to encode queries and documents into dense vector spaces. The relevance of documents is assessed based on the similarity of these vectors, allowing to capture of semantic relationships that go beyond keyword matching. Regardless of the retrieval technique employed, the retrieval system identifies the top-K most relevant documents, denoted as  $\mathcal{D} = \{d_1, \dots, d_K\}$ .

### 2.2 Answer Scent Generation

Large language models (LLMs) such as GPT-3 (Brown et al., 2020) and LLaMA (Touvron et al., 2023) define probability distributions over sequences of tokens. Given a sequence  $x_1, \dots, x_n$ , these models typically predict the sequence’s probability using an autoregressive approach  $p(x_1, \dots, x_n) = \prod_{i=1}^n p(x_i|x_{<i})$ , where  $x_{<i} := x_1, \dots, x_{i-1}$  represents the sequence of tokens preceding  $x_i$ , also referred to as its *prefix*. This modeling is implemented via a transformer network parameterized by  $\theta_1$ , typically employing a causal self-attention mask  $p(x_1, \dots, x_n) = \prod_{i=1}^n p_{\theta_1}(x_i|x_{<i})$ , which effectively models the conditional probabilities of each token.

In our approach, we incorporate the concept of *answer scent*, which guides the model in generating answers that are contextually appropriate for the query, inspired by the success of In-Context Learning (Brown et al., 2020; Ram et al., 2023; Dong et al., 2022). This context is derived using a zero-shot approach, where the model infers the scent without explicit prior training on such task:  $p(x_1, \dots, x_n) = \prod_{i=1}^n p(x_i|x_{<i}, \mathcal{S}(x_{<i}))$ , where  $\mathcal{S}(x_{<i})$  denotes the inferred answer scent in the form of a generated text content. The objective of Scent is to encode the essence of what the answer should represent, enriching the input provided to the LLM reranker.

### 2.3 ASRANK Re-Ranking

ASRANK introduces an unsupervised re-ranking utilizing LLM to evaluate the relevance of documents based on  $\mathcal{S}(q)$ , which serves as the guiding context corresponding to the target query  $q$ .

The core of ASRANK’s method is the calculation of a relevance score for each document, leveraging

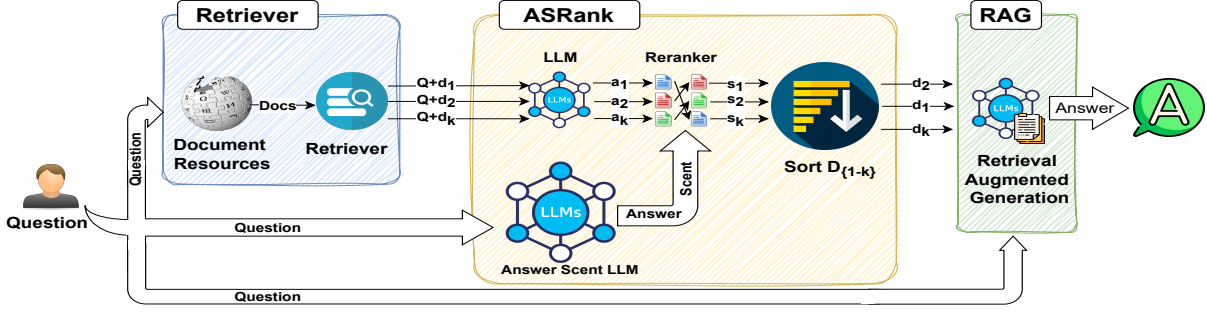


Figure 2: Our ASRANK framework, starts with document retrieval, re-ranking using the answer scent from LLMs, and finally passing the top-k document into the RAG system.

both the content of the document and its alignment with the inferred answer scent. The score is formulated as:

$$s(\mathbf{d}_i) = \sum_{t=1}^{|\mathbf{a}|} -\log p(a_t | \mathbf{a}_{<t}, \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q}); \theta_2),$$

where  $\mathbf{d}_i$  represents an individual document within the set of retrieved documents  $\mathcal{D}$ ,  $|\mathbf{a}|$  denotes the length of the  $\mathbf{a}$  - an answer generated based on  $\mathbf{d}_i$  by the rank model, and  $\mathcal{S}(\mathbf{q})$  represents the answer scent derived from the query  $\mathbf{q}$ . The term  $\log p(a_t | \mathbf{a}_{<t}, \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q}); \theta_2)$  is the log probability of each token  $a_t$  of the answer conditional on the prior tokens  $\mathbf{a}_{<t}$ , the document  $\mathbf{d}_i$ , the query  $\mathbf{q}$ , and the answer scent, parameterized by the model’s parameters  $\theta_2$ .

To elaborate, the computation of the conditional probabilities can be decomposed as follows:

$$\log p(\mathbf{a} | \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q})) = \sum_{t=1}^{|\mathbf{a}|} \log p(a_t | \mathbf{a}_{<t}, \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q}); \theta_2),$$

where  $\log p(\mathbf{a} | \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q}))$  represents the log probability of generating the answer  $\mathbf{a}$  given the document  $\mathbf{d}_i$ , the query  $\mathbf{q}$ , and the inferred answer scent  $\mathcal{S}(\mathbf{q})$ . The relevance score, denoted by  $s(\mathbf{d}_i)$ , is reformulated using Bayes’ Theorem. This score is represented as follows:

$$s(\mathbf{d}_i) \propto \log p(\mathbf{a} | \mathbf{d}_i, \mathbf{q}, \mathcal{S}(\mathbf{q})) + \log p(\mathbf{d}_i | \mathbf{q}) - \log p(\mathbf{a} | \mathbf{q}),$$

where  $\log p(\mathbf{d}_i | \mathbf{q})$  is the log probability that the document  $\mathbf{d}_i$  is relevant to the query  $\mathbf{q}$ , based on the initial retrieval. The normalization term  $-\log p(\mathbf{a} | \mathbf{q})$  adjusts for the base likelihood of the answer  $\mathbf{a}$  being related to the query  $\mathbf{q}$  across all documents.

The decision to select the most relevant document employs a maximization approach  $\hat{i} =$

$\arg \max_{i \in [1, K]} s(\mathbf{d}_i)$ , enhancing the likelihood that the document contains the information necessary to answer the query effectively, aligned with the derived answer scent.

## 3 Experiment Settings

### 3.1 Datasets

We utilize several common datasets for our experiments, whose detailed statistics are also provided in Appendix A:

**Open Domain QA datasets:** TriviaQA (Joshi et al., 2017) is a collection of trivia questions sourced from trivia and quiz-league websites. Natural Questions (NQ) (Kwiatkowski et al., 2019) is a question-answering dataset containing 79,168 training examples, 8,757 development examples, and 3,610 test question-answer pairs. WebQuestions (Berant et al., 2013) is a question-answering dataset that was created using Freebase as a knowledge base and which contains 5,810 question-answer pairs.

**Entity-centric Questions:** EntityQuestions (Scialvolino et al., 2021) contains 22K short questions about named entities based on facts from Wikipedia.

**Temporal Questions:** ArchivalQA (Wang et al., 2022a) is a large-scale question answer collection designed specifically for temporal news QA, containing 532,444 question-answer pairs, often on detailed or minor aspects. These pairs are derived from the New York Times Annotated Corpus, which spans news articles published between January 1, 1987, and June 19, 2007. We follow prior work (Wallat et al., 2024) and evaluate ASRANK on the subset of ArchivalQA dataset, which comprises 7,500 questions.

**Multi-hop Questions:** HotpotQA (Yang et al., 2018) contains 113K crowd-sourced questions con-

structured in a way that the introduction paragraphs of two Wikipedia articles are required to answer questions (i.e., two hops). We focus on the full-wiki setting, in which two Wikipedia passages are required to answer the questions. We follow prior work (Khalifa et al., 2022) and evaluate ASRANK on the development set, which has 7,405 questions.

**Information Retrieval:** BEIR is a diverse suite designed to test retrieval algorithms across multiple tasks, such as fact-checking and question-answering, with datasets from various domains including news, technical documents, and Wikipedia (Thakur et al., 2021). Specifically, we utilize the NFCorpus, DBPedia, Touche and News datasets from BEIR, which represent a range of retrieval challenges. TREC-DL19 (Craswell et al., 2020) is a benchmark dataset widely used in IR research. We use the test sets of 2019 which contain 43 queries.

### 3.2 Retrievers

In our re-ranking experiments, we first retrieve passages using both unsupervised and supervised retrievers, as detailed below.

**Unsupervised Retrievers:** BM25 (Robertson et al., 2009) is a ranking function used by search engines to estimate the relevance of documents to a given search query. Masked Salient Spans (MSS) (Sachan et al., 2021a) is a dense retriever trained by predicting masked salient spans like named entities with the help of a reader network. Contriever (Izacard et al., 2022) is a framework for pre-training and fine-tuning models for information retrieval using contrastive learning.

**Supervised Retrievers:** Dense Passage Retrieval (DPR) (Karpukhin et al., 2020) uses annotated question-context paragraphs and hard negative examples to train a supervised dense retriever. MSS-DPR (Sachan et al., 2021a) further improves DPR performance by first pre-training the dense retriever using MSS followed by DPR-style supervised fine-tuning.

### 3.3 LLM Models

This section overviews the large language models (LLMs) utilized in our experiments. These models are essential for generating the "answer scent" and for re-ranking documents based on their inferred relevance to the query.

**Answer Scent Models:** We leverage a variety of Large Language Models (LLMs) to generate the answer sent, each bringing unique strengths to our re-ranking method. The **Llama** models, developed by Meta, are known for their robust performance in dialogue applications, having undergone extensive pre-training and fine-tuning (Touvron et al., 2023). **Mistral** and **Mixtral**, from Mistral AI, push the boundaries of efficiency and computational optimization, employing instruction fine-tuning and a sparse mixture of experts approach, respectively (Jiang et al., 2023, 2024). **Gemma**, a product of Google, offers both base and instruction-tuned versions in different sizes, designed for adaptability across various hardware platforms (Team et al., 2024). **GPT**, from OpenAI, is renowned for its general-purpose capabilities, pre-trained on vast data pools to generate semantically rich responses (Brown et al., 2020). Lastly, **Qwen**, by Alibaba Cloud, encapsulates a broad pre-training regime across multiple languages and domains, optimized for long-context interactions, highlighting its scalability and depth in handling complex linguistic tasks (Bai et al., 2023).

**Rank Model:** In our experiments, we specifically utilize the T5 Base and T5 Large models, a variation of the original T5 architecture (Raffel et al., 2020) adapted for language modelling tasks. This architecture features encoder and decoder transformers pre-trained to improve their ability to handle input text sequences.

### 3.4 Experimental Setup

All re-ranking experiments were conducted on a high-performance computing cluster using NVIDIA A100 48GB GPUs, while some experiments, such as ones in Section 6 were done using NVIDIA A40 GPUs. All experiments and dataset processing were conducted using the Rankify framework, which provides a unified toolkit for retrieval, re-ranking, and retrieval-augmented generation (Abdallah et al., 2025b). We evaluated our method across five retrievers: BM25, MSS, MSS-DPR, DPR, and Contriever, retrieving the top 1,000 passages for re-ranking, consistent with the setup in Sachan et al. (2022). For temporal questions in the ArchivalQA dataset, we also included Ance (Xiong et al., 2020) and RocketQA (Qu et al., 2020) for a comprehensive evaluation. Retrieval settings for the HotpotQA dataset followed the configurations from Khalifa et al. (2022) to en-

Retriever	NQ				TriviaQA				WebQ			
	Top-1	Top-5	Top-10	Avg	Top-1	Top-5	Top-10	Avg	Top-1	Top-5	Top-10	Avg
<i>Unsupervised Retrievers</i>												
MSS	19.2	41.2	51.2	37.2	30.7	52.6	60.5	47.9	11.6	29.0	39.1	26.6
MSS + UPR <sub>T0_3B</sub>	38.7	<b>64.8</b>	<b>72.2</b>	58.6	57.2	75.5	78.9	70.5	29.9	57.4	65.0	57.7
MSS + ASRANK †+T5 <sub>base</sub>	45.2	64.7	70.6	60.1	65.3	77.2	79.8	74.1	42.5	61.3	67.7	50.1
MSS + ASRANK ‡+T5 <sub>base</sub>	<b>46.5</b>	64.4	69.8	<b>60.2</b>	<b>66.3</b>	<b>77.6</b>	<b>80.1</b>	<b>74.6</b>	<b>45.0</b>	<b>63.6</b>	<b>68.8</b>	<b>59.1</b>
BM25	22.1	43.7	54.4	40.1	46.3	66.2	71.7	61.4	18.8	41.8	52.1	37.6
BM25 + UPR <sub>T0_3B</sub>	35.4	63.4	70.2	56.3	55.7	76.5	80.2	70.8	30.1	57.3	66.5	51.3
BM25 + RANKGPT ‡	43.4	62.3	68.0	57.9	-	-	-	-	40.3	57.6	64.1	54.0
BM25 + ASRANK †+T5 <sub>base</sub>	46.2	65.3	<b>72.3</b>	61.2	67.2	<b>77.9</b>	<b>80.7</b>	75.2	44.8	63.7	68.7	<b>59.0</b>
BM25 + ASRANK ‡+T5 <sub>base</sub>	<b>47.3</b>	<b>65.6</b>	71.4	<b>61.4</b>	<b>67.3</b>	<b>77.9</b>	<b>80.7</b>	<b>75.3</b>	<b>45.4</b>	<b>62.9</b>	<b>68.9</b>	<b>59.0</b>
Contriever	22.1	47.2	58.7	42.7	34.1	59.4	68.0	53.8	19.9	43.4	56.3	39.9
Contriever + UPR <sub>T0_3B</sub>	36.4	64.6	72.4	57.8	56.7	77.0	80.2	71.3	30.0	58.5	68.2	52.2
Contriever + RANKGPT ‡	44.4	64.8	69.6	59.6	-	-	-	-	43.6	63.6	70.0	59.0
Contriever + ASRANK †+T5 <sub>base</sub>	41.5	61.3	68.4	57.0	57.9	72.8	76.8	69.1	42.9	62.7	69.8	58.4
Contriever + ASRANK ‡+T5 <sub>base</sub>	<b>48.0</b>	<b>66.6</b>	<b>72.5</b>	<b>62.3</b>	<b>66.8</b>	<b>78.9</b>	<b>81.4</b>	<b>76.0</b>	<b>46.8</b>	<b>64.8</b>	<b>70.8</b>	<b>60.8</b>
<i>Supervised Retrievers</i>												
DPR	48.6	68.7	74.5	63.9	57.4	72.4	76.5	68.7	44.8	65.0	70.6	60.1
DPR + UPR <sub>T0_3B</sub>	42.5	70.6	<b>78.1</b>	63.8	61.3	78.7	81.9	74.0	34.9	63.6	71.7	56.7
DPR + RANKGPT ‡	48.6	68.7	74.5	63.9	-	-	-	-	44.8	65.0	70.6	60.1
DPR + ASRANK †+T5 <sub>base</sub>	50.2	69.9	76.1	65.3	68.8	79.8	<b>82.4</b>	77.0	48.2	<b>68.1</b>	73.2	63.1
DPR + ASRANK ‡+T5 <sub>base</sub>	<b>51.3</b>	<b>70.6</b>	76.0	<b>65.9</b>	<b>69.9</b>	<b>79.8</b>	82.1	<b>77.3</b>	<b>49.3</b>	67.3	<b>73.4</b>	<b>63.3</b>
MSS-DPR	50.1	<b>71.8</b>	<b>77.4</b>	<b>66.5</b>	61.6	75.2	79.1	71.9	44.2	65.0	71.6	60.3
MSS-DPR + UPR <sub>T0_3B</sub>	41.4	69.8	77.9	63.0	60.5	78.9	82.5	74.0	31.8	61.6	70.3	54.5
MSS-DPR + ASRANK †+T5 <sub>base</sub>	48.8	69.3	76.1	64.7	69.4	80.4	82.9	77.5	47.7	<b>67.0</b>	<b>73.0</b>	62.5
MSS-DPR + ASRANK ‡+T5 <sub>base</sub>	<b>50.6</b>	69.3	75.2	65.0	<b>69.9</b>	<b>80.5</b>	<b>82.9</b>	<b>77.7</b>	<b>49.7</b>	66.6	72.6	<b>62.9</b>

Table 1: Top-1, 5, 10 retrieval accuracy of re-ranking methods including ASRANK and baseline models on the NQ, TriviaQA and WebQ Datasets. † refers to Llama-3-70B, ‡ refers to GPT-3.5-turbo-0125. For a comparison between LLama-2 7b vs UPR see Table 12 in Appendix E.8. Due to the Computational cost of RankGPT, we experimented with three retrievers (BM25, contriever, DPR) on two datasets (NQ, WebQA).

sure consistency. The UPR baseline was primarily evaluated using the T0-3B model, except in Table 3, where T0-3B, T0-11B, and T5-11B were used to explore model size variations. For BEIR and TREC datasets, we adopted the same settings as RankGPT (Sun et al., 2023), using BM25 with 100 retrieved documents. To generate the answer scent, we employed the Llama-3-70B (instruction version) and GPT-3.5-turbo-0125 models, which were set to a temperature of 0.7 and a max length of 128 tokens. For re-ranking, T5-base and T5-large models were used, with a batch size of 128. To evaluate ASRANK’s performance, we used Top-K retrieval accuracy, following the methodology outlined in Sachan et al. (2022). For RAG evaluations, we measured performance using exact match, recall, and F1 scores. Additional details on the framework implementation and metrics can be found in Appendix B and D.

## 4 Experiment Results

In this section, we evaluate ASRANK on a variety of question-answering tasks, leveraging several datasets to assess its performance. The datasets employed cover different QA challenges, ranging from open domain to entity-centric, temporal, BEIR, TREC and multi-hop questions. The primary objective is to evaluate ASRANK’s capability to rank the Top-1, 5, 10 retrieved passages. For this purpose, an initial retrieval of 1,000 passages per question is conducted for reranking. On

benchmarks, we compare ASRANK with LLama 70B and GPT3.5-Turbo with state-of-the-art supervised and unsupervised passage re-ranking methods. The baselines include UPR (Sachan et al., 2022), RankGPT (Sun et al., 2023), HYDE (Gao et al., 2022), PathRetriever (Asai et al., 2020), MDR (Xiong et al., 2021), DrKit (Dhingra et al., 2020), and PromptRank (Khalifa et al., 2022).

### 4.1 ODQA Re-ranking

In this section, we focus on evaluating ASRANK across several ODQA datasets (NQ, TriviaQA, and WebQ). Table 1 shows improvements in retrieval Top-K accuracy. ASRANK enhances the retrieval of Top-K results across various settings, often outperforming the UPR model. For instance, when combined with the MSS retriever on NQ dataset, ASRANK † increases the relevance of the Top-1 result to 45.2%, a notable improvement over UPR’s increase to 38.7%. Similarly, TriviaQA dataset, ASRANK with BM25 achieves a Top-1 accuracy of 47.3%, surpassing BM25 + UPR’s performance of 35.4%. Also, the combination of ASRANK with the MSS retriever results in a remarkable uplift in Top-1 accuracy for NQ, from an initial 19.2% to 46.5%. Similarly, on TriviaQA, ASRANK with the BM25 retriever increases the accuracy of Top-1 from 22.1% to 47.3%. In Appendix G, we show random examples from NQ dev and WebQA after and before re-ranking. Further analysis of the effect of answer scent lengths is presented in Ap-

pendix E.4. The performance improvements across various LLMs are detailed in Appendix E.5

Additionally, we compare ASRANK with RankGPT in Table 1. ASRANK outperforms RankGPT in retrieval accuracy across the NQ and WebQ datasets. For instance, ASRANK ‡ (i.e., ASRANK with Llama-3-70B) achieves a Top-1 accuracy of 47.3% on NQ, compared to RankGPT’s 43.4% with the BM25 retriever. In WebQ, ASRANK ‡ reaches a Top-1 accuracy of 45.4%, surpassing RankGPT’s 40.3%. **Cost:** Utilizing RankGPT with GPT-3.5 across NQ and WebQ datasets, with DPR, Contriever, and BM25 retrievers, incurs a total cost of \$700. In comparison, running the same experiments with ASRANK costs only \$15, demonstrating its cost-effectiveness at scale. Further details on performance improvements can be found in Appendix E.1 while the impact of answer scent on re-ranking and latency implications are reported in Appendix E.6.

Method	Model	DL 19	NFCorpus	Touche	DBPedia	News	Avg
BM25	-	50.58	30.75	<b>44.22</b>	31.80	39.52	39.37
UPR	T5 <sub>VL</sub>	53.85	35.04	19.69	30.91	43.11	36.52
monoBERT	BERT	70.50	36.88	31.75	41.87	44.62	45.12
monoT5	T5 <sub>Base</sub>	71.48	37.38	30.82	42.42	46.83	45.78
Cohere Rerank-v2	-	73.22	36.36	32.51	42.51	47.59	46.43
Promptagator++ (few-shot)	-	-	37.0	38.10	43.4	-	-
RankGPT	GPT-3.5	65.80	35.62	36.18	44.47	<b>48.85</b>	46.18
HyDE	InstructGPT	61.30	-	-	36.80	44.00	-
ASRANK	Llama-3-70B+T5 <sub>Base</sub>	72.15	33.90	33.22	43.51	41.13	44.78
ASRANK	GPT-3.5+T5 <sub>Base</sub>	<b>74.42</b>	<b>38.15</b>	35.56	<b>45.37</b>	48.45	<b>48.39</b>

Table 2: Results (nDCG@10) on TREC and BEIR. GPT-3.5 is turbo-0125.

## 4.2 BEIR and TREC Results

In this section, we evaluate the performance of ASRANK on both the TREC (DL 19) and BEIR datasets, using nDCG@10 as the evaluation metric. For each dataset, the retriever selected the top 100 documents using BM25. Table 2 presents the results from the DL 19 and BEIR datasets, which includes NFCorpus, DBPedia, and News. Across the datasets, ASRANK method, using GPT-3.5 combined with T5 base, outperformed the other approaches, achieving the highest average score of 48.58. Specifically, ASRANK achieved a score of 74.42 on the DL 19 dataset, surpassing both RankGPT and HyDE. While RankGPT performed slightly better on the News dataset, reaching 48.85 compared to ASRANK’s 48.45, ASRANK remained competitive across all datasets. The Llama 70B + T5 base variant of ASRANK also delivered strong results, particularly on the DL 19 dataset, scoring 72.15.

## 5 Additional Studies

### 5.1 Evaluation on NQ development set

In this section, we compare our approach with UPR using different model sizes (T0-3B, T0-11B) (Sanh et al., 2021) to assess the efficacy in the context of the NQ development set. This comparison highlights the significant advantages offered by the ASRANK, across different retrievers like BM25, MSS, and DPR. The results are shown in Table 3. ASRANK enhances retrieval performance across all Top-1, 5, 10, 20 results. Notably, after re-ranking with ASRANK using the Llama 70B configuration, the accuracy of Top-1 for the MSS-DPR combination reaches 48.1%, which is an improvement over its performance with UPR, where the Top-1 achieves 39.7%. Similarly, the Top-1 for DPR alone ascends to 50.4% with GPT3.5, surpassing the 41.1% recorded with UPR. Additionally, the performance of ASRANK with T5-base + BM25 achieves a Top-1 accuracy of 47.5%, surpassing previous results obtained with T5-lm-adapt (11B) or T0-11B.

### 5.2 Evaluation on Diverse Question Answering Datasets

The further studies were conducted across three distinct datasets—EntityQuestions, HotpotQA, and ArchivalQA. As summarized in Tables 4, 7 (Appendix E.2), and 8 (Appendix E.3), ASRANK enhances Top-1, 5, and 10 retrieval accuracies, across different retrievers.

The EntityQuestions dataset, when used Llama 70B and GPT3.5 boosts performance, achieving Top-1 retrieval accuracy up to 48.9%, which indicates an increase of over 25% compared to baselines. For the HotpotQA dataset, which requires reasoning over multiple documents, ASRANK achieves substantial enhancements in Top-2, surpassing fully-supervised baselines like DPR (Karpukhin et al., 2020) and DrKit (Dhingra et al., 2020), MDR (Xiong et al., 2021), PathRetriever (Asai et al., 2020) when combined with TF-IDF. This shows ASRANK’s strength in multi-hop question answering, supporting complex inference tasks across linked data points. Notably, ASRANK combined with DPR achieves a Top-2 accuracy of 42.6%, which not only surpasses the fully-supervised baselines such as DPR at 18.5% and DrKit at 38.3%, but also outperforms unsupervised models like PromptRank-GPT2-XL and PromptRank-T5-XL (Khalifa et al., 2022), which

Retriever	NQ (dev)				Avg
	Top-1	Top-5	Top-10	Top-20	
MSS	17.7	38.6	48.7	57.4	40.6
BM25	22.3	43.8	53.6	62.3	45.5
Contriever	19.6	45.4	55.8	64.9	46.4
MSS+BM25	17.6	38.7	48.8	57.8	40.7
MSS-DPR	<b>48.9</b>	<b>69.9</b>	<b>75.7</b>	<b>80.4</b>	<b>68.7</b>
DPR	47.8	67.3	73.0	77.4	6.4
After Re-ranking with UPR					
MSS+T0-3B	36.6	62.9	70.8	75.7	61.5
BM25+MSS+T5-lm-adapt (3B)	29.7	59.9	-	76.9	55.5
BM25+MSS+T5-lm-adapt (11B)	32.1	62.3	-	78.5	57.6
BM25+MSS+ T0-3B	36.7	64.9	-	79.1	60.2
BM25+MSS+ T0-11B	37.4	64.9	-	79.1	60.5
MSS+DPR+T0-3B	39.7	68.6	76.5	<b>82.0</b>	66.7
DPR+T0-3B	<b>41.1</b>	<b>69.5</b>	<b>77.0</b>	81.9	<b>67.4</b>
After Re-ranking with Llama-3-8B					
BM25+T5 <sub>base</sub>	39.8	58.6	66.6	71.9	59.6
BM25+MSS+T5 <sub>Large</sub>	<b>41.7</b>	<b>64.8</b>	<b>73.5</b>	<b>80.1</b>	<b>65.0</b>
After Re-ranking with GPT-3.5-turbo-0125					
MSS+T5 <sub>base</sub>	46.2	63.5	69.1	73.2	63.0
BM25+T5 <sub>base</sub>	47.5	64.5	69.7	74.3	64.0
Contriever+T5 <sub>base</sub>	47.7	65.5	71.2	76.2	65.2
BM25+MSS+T5 <sub>base</sub>	47.9	65.5	71.2	76.4	65.3
MSS-DPR+T5 <sub>base</sub>	50.1	68.9	74.8	<b>79.8</b>	68.4
DPR+T5 <sub>base</sub>	<b>50.4</b>	<b>68.9</b>	<b>74.9</b>	79.4	<b>68.4</b>
After Re-ranking with Llama-3-70B					
MSS+T5 <sub>base</sub>	44.9	63.7	69.4	73.9	62.9
BM25+T5 <sub>base</sub>	44.8	64.1	69.9	75.0	63.5
Contriever+T5 <sub>base</sub>	45.7	65.4	71.5	76.2	64.7
BM25+MSS+T5 <sub>base</sub>	45.4	65.4	71.0	76.6	64.6
MSS-DPR+T5 <sub>base</sub>	48.1	<b>68.6</b>	<b>74.6</b>	<b>79.8</b>	<b>67.8</b>
DPR+T5 <sub>base</sub>	<b>48.2</b>	67.9	73.8	78.6	67.1

Table 3: Performance comparison of different retrievers on the NQ development set.

score 36.6% and 42.8%, respectively. On the ArchivalQA dataset, which contains temporal questions, ASRANK shows also good improvements. After re-ranking with Llama 70B and GPT3.5, the model significantly boosts Top-1 accuracies across different retrievers, demonstrating its effectiveness in extracting temporally relevant information. Specifically, after re-ranking with Llama 70B, BM25 improves from 18.2% to 26.2% in Top-1, DPR from 17.0% to 27.5%, and ANCE from 18.0% to 27.3%. When using GPT3.5, BM25 improves further to 27.6% in Top-1, DPR reaches 27.7%, and ANCE advances to 28.1%.

Retriever	HotpotQA			
	# Ex.	top-2	top-10	top-20
Fully-supervised Baselines				
DPR	-	18.5	37.2	47.1
DPR+ASRANK	-	42.6	68.8	79.2
DrKit	-90K	38.3	67.2	71.0
MDR	-90K	65.9	77.5	<b>80.2</b>
PathRetriever	-90K	<b>66.4</b>	<b>77.8</b>	78.7
Unsupervised Baselines				
TF-IDF	-	9.9	27.6	35.0
TF-IDF + BM25	-	19.1	54.7	61.8
PromptRank-GPT2-XL	-	36.6	60.5	65.9
PromptRank-T5-XL	-	42.8	68.9	74.1
TF-IDF+ASRANK †+T5 <sub>base</sub>	-	36.9	61.1	72.5
TF-IDF+ASRANK ‡+T5 <sub>base</sub>	-	<b>45.1</b>	<b>69.1</b>	<b>78.9</b>

Table 4: Top-2, 10, 20 retrieval performance on HotpotQA comparing ASRANK to baselines. † refers to Llama-3-70B and ‡ refers to GPT-3.5-turbo-0125.

### 5.3 Role of Answer Scent

In this section, we present an ablation study evaluating the contribution of the answer scent in the

ASRANK framework. Table 5 shows the results for different retrievers (DPR, MSS, and BM25) under various conditions: with an unknown (<UNK>) token used as the answer scent, and with the answer scent on NQ-test. For instance, DPR’s Top-1 accuracy improves from 23.4% without answer scent to 51.3% when the answer scent is provided. Similarly, BM25’s Top-1 accuracy increases from 23.5% to 47.3%.

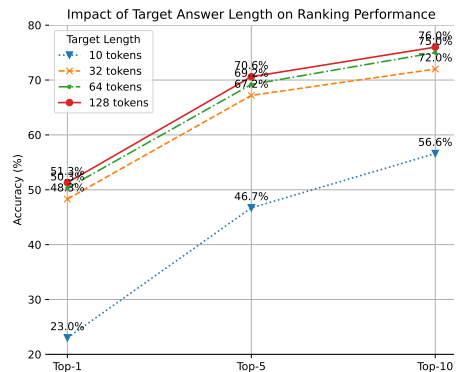


Figure 3: Impact of target answer length on ranking performance using DPR on NQ-test.

Retriever	Mode	Top-1	Top-5	Top-10
MSS	<UNK>	22.9	29.78	46.9
	Answer	46.5	64.4	69.8
BM25	<UNK>	23.5	30.6	45.9
	Answer	47.3	65.6	71.4
DPR	<UNK>	23.4	31.9	49.4
	Answer	51.3	70.6	76.0

Table 5: Performance Comparison with and without Answer Scent for NQ test dataset.

### 5.4 Impact of Answer Length

We conducted two sets of experiments to further analyze the potential influence of answer length (Wang et al., 2022b) on ranking performance. The first experiment focused on the length of the generated answer scent, while the second examined the target length of answers provided by the T5 model. Table 9 presents the results for different lengths of the answer scent. We tested varying lengths (50, 75, 100, and 128 tokens) and found that 128 tokens consistently produced better results regarding Top-1, Top-5, and Top-10 accuracy, with DPR’s Top-1 accuracy reaching 51.3%. Fig 3 explores the impact of different target answer lengths on ranking accuracy. The results show that increasing the target answer length generally improves the

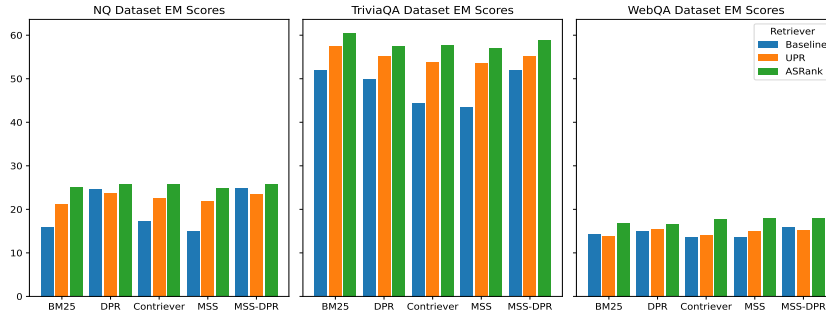


Figure 4: Comparison of Exact Match (EM) scores across three datasets (NQ, TriviaQA, and WebQA) for various retrieval models.

ranking, with the best performance observed at a target length of 128 tokens.

## 6 RAG for ODQA

**Method** In the Retrieval-Augmented Generation (RAG) framework, we employ a large language model (LLM), leveraging its capacity to utilize retrieved documents dynamically for generating responses. The RAG method combines the robust retrieval capabilities of DPR with the generative models, thereby enabling understanding and response generation based on the context provided by the retrieved documents. The RAG model is formulated as  $p(a | q, D) = \sum_{d \in D} p(d | q) \cdot p(a | q, d)$ , where  $a$  represents the answer,  $q$  the query, and  $D$  is the set of retrieved documents relevant to  $q$ . The term  $p(d | q)$  denotes the document’s retrieval probability, and  $p(a | q, d)$  represents the probability of generating answer  $a$  given the query  $q$  and document  $d$ .

**Results** We evaluated the RAG method on NQ, TriviaQA, and WebQA revealing significant performance gains as shown in Figure 4. For instance, before applying our ASRANK re-ranking strategy, the BM25+LLama7B achieves a baseline EM of 16.0% on NQ. After re-ranking with ASRANK, the EM increased to 24.8%. We show a detailed comparison between the baselines (BM25, DPR, MSS, Contriever, MSS-DPR), UPR, and ASRANK in Table 13 (Appendix F).

## 7 Related Work

Recent developments in information retrieval have focused on integrating LLMs to enhance retrieval and reranking mechanisms. Approaches like InPars (Bonifacio et al., 2022), Promptagator (Dai et al., 2022), DynRank (Abdallah et al., 2025a) and UPR (Sachan et al., 2022) demonstrate how syn-

thetic data generation and LLM-based scoring improve domain-specific retrieval performance. Models like SGPT (Muennighoff, 2022) and PRP (Qin et al., 2023) have highlighted the impact of fine-tuning LLMs such as LLaMA for reranking, with PRP significantly enhancing retrieval compared to smaller models. Traditional methods like BM25 (Robertson et al., 2009), MSS (Sachan et al., 2021a), and DPR (Karpukhin et al., 2020) are pivotal in initial retrieval stages, further refined by dense and sparse techniques like MSS-DPR (Sachan et al., 2021b). Our work is related to (Izacard and Grave, 2021), which transfers knowledge from a reader model to a retriever, paralleling our use of large models to enhance retrieval quality via answer scent generation. Similarly, recent efforts such as RankVicuna (Pradeep et al., 2023) and LRL (Ma et al., 2023) demonstrate how LLMs efficiently reorder documents using prompts, reflecting the trend towards prompt-based reranking with LLMs.

## 8 Conclusions

In this paper, we introduced ASRANK, a novel zero-shot re-ranking method that leverages the concept of answer scent to improve document retrieval for open-domain question answering and information retrieval tasks. Through comprehensive experiments on diverse datasets, including both open-domain and document-centric benchmarks like BEIR and TREC, we demonstrated that ASRANK consistently outperforms unsupervised and supervised baselines. Our approach made significant improvements in Top-1, Top-5, and Top-10 retrieval accuracy, particularly when integrated with LLMs like GPT-3.5 and Llama-70B. Moreover, the cost-effectiveness and computational efficiency of ASRANK, compared to other large-scale reranking methods, underscore its practical utility.



## Limitations

While ASRANK demonstrates significant improvements in document re-ranking with the incorporation of answer scent, there are several limitations that warrant discussion:

1. The computational cost associated with ASRANK increases with the number of documents due to the need to compute the score based on the answer scent with the answer generated from each document.
2. The effectiveness and consistency of ASRANK are contingent upon the specific pre-trained language models used for generating the answer scent. Variations in these models, due to different training data or updates, can introduce biases and affect the stability of the re-ranking outcomes.
3. ASRANK's performance heavily depends on the quality of the initial retrieval phase.

## Ethical Considerations and Licensing

Our research utilizes the GPT models, which is available under the OpenAI License and Apache-2.0 license, and the Llama model, distributed under the Llama 2 Community License Agreement provided by Meta. We ensure all use cases are compliant with these licenses. Additionally, the datasets employed are sourced from repositories permitting academic use. We are releasing the artifacts developed during our study under the MIT license to facilitate ease of use and adaptations by the research community. We have ensured that all data handling, model training, and dissemination of results are conducted in accordance with ethical guidelines and legal stipulations associated with each used artifact.

## References

- Abdelrahman Abdallah, Jamshid Mozafari, Bhawna Piryani, Mohammed M. Abdelgwad, and Adam Jatowt. 2025a. [DynRank: Improve passage retrieval with dynamic zero-shot prompting based on question classification](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4768–4778, Abu Dhabi, UAE. Association for Computational Linguistics.
- Abdelrahman Abdallah, Bhawna Piryani, Jamshid Mozafari, Mohammed Ali, and Adam Jatowt. 2025b.

[Rankify: A comprehensive python toolkit for retrieval, re-ranking, and retrieval-augmented generation](#). Preprint, arXiv:2502.02464.

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. [Learning to retrieve reasoning paths over wikipedia graph for question answering](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*.
- Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. Inpars: Unsupervised dataset generation for information retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2387–2392.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the trec 2019 deep learning track. *arXiv preprint arXiv:2003.07820*.
- Zhuyun Dai, Vincent Y Zhao, Ji Ma, Yi Luan, Jianmo Ni, Jing Lu, Anton Bakalov, Kelvin Guu, Keith B Hall, and Ming-Wei Chang. 2022. [Promptagator: Few-shot dense retrieval from 8 examples](#). *arXiv preprint arXiv:2209.11755*.
- Bhuwan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig, Ruslan Salakhutdinov, and William W. Cohen. 2020. [Differentiable reasoning over a virtual knowledge base](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. [A survey on in-context learning](#). *arXiv preprint arXiv:2301.00234*.
- Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2022. [Precise zero-shot dense retrieval without relevance labels](#). *arXiv preprint arXiv:2212.10496*.

- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Transactions on Machine Learning Research*.
- Gautier Izacard and Edouard Grave. 2021. [Distilling knowledge from reader to retriever for question answering](#). In *International Conference on Learning Representations*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Muhammad Khalifa, Lajanugen Logeswaran, Moon-tae Lee, Honglak Lee, and Lu Wang. 2022. Few-shot reranking for multi-hop qa via language model prompting. *arXiv preprint arXiv:2205.12650*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: a benchmark for question answering research](#). *Transactions of the Association of Computational Linguistics*.
- Jakub Lála, Odhran O’Donoghue, Aleksandar Shtedritski, Sam Cox, Samuel G Rodrigues, and Andrew D White. 2023. Paperqa: Retrieval-augmented generative agent for scientific research. *arXiv preprint arXiv:2312.07559*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Xueguang Ma, Xinyu Zhang, Ronak Pradeep, and Jimmy Lin. 2023. Zero-shot listwise document reranking with a large language model. *arXiv preprint arXiv:2305.02156*.
- David Maxwell and Leif Azzopardi. 2018. Information scent, searching and stopping: Modelling serp level stopping behaviour. In *Advances in Information Retrieval: 40th European Conference on IR Research, ECIR 2018, Grenoble, France, March 26-29, 2018, Proceedings 40*, pages 210–222. Springer.
- Niklas Muennighoff. 2022. Sgpt: Gpt sentence embeddings for semantic search. *arXiv preprint arXiv:2202.08904*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems*.
- Ronak Pradeep, Sahel Sharifymoghaddam, and Jimmy Lin. 2023. Rankvicuna: Zero-shot listwise document reranking with open-source large language models. *arXiv preprint arXiv:2309.15088*.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, et al. 2023. Large language models are effective text rankers with pairwise ranking prompting. *arXiv preprint arXiv:2306.17563*.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2020. Rocketqa: An optimized training approach to dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2010.08191*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Devendra Singh Sachan, Mike Lewis, Mandar Joshi, Armen Aghajanyan, Wen-tau Yih, Joelle Pineau, and

- Luke Zettlemoyer. 2022. Improving passage retrieval with zero-shot question generation. *arXiv preprint arXiv:2204.07496*.
- Devendra Singh Sachan, Mostofa Patwary, Mohammad Shoeybi, Neel Kant, Wei Ping, William L Hamilton, and Bryan Catanzaro. 2021a. [End-to-end training of neural retrievers for open-domain question answering](#). In *Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Devendra Singh Sachan, Siva Reddy, William L. Hamilton, Chris Dyer, and Dani Yogatama. 2021b. [End-to-end training of multi-document reader and retriever for open-domain question answering](#). In *Advances in Neural Information Processing Systems*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. 2021. [Multi-task prompted training enables zero-shot task generalization](#). *Preprint*, arXiv:2110.08207.
- Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. [Simple entity-centric questions challenge dense retrievers](#). In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Spurthi Setty, Katherine Jijo, Eden Chung, and Natan Vidra. 2024. Improving retrieval for rag based question answering models on financial documents. *arXiv preprint arXiv:2404.07221*.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is chatgpt good at search? investigating large language models as re-ranking agents](#). *Preprint*, arXiv:2304.09542.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jonas Wallat, Adam Jatowt, and Avishek Anand. 2024. Temporal blind spots in large language models. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 683–692.
- Jiexin Wang, Adam Jatowt, and Masatoshi Yoshikawa. 2022a. Archivalqa: a large-scale benchmark dataset for open-domain question answering over historical news collections. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3025–3035.
- Yequan Wang, Jiawen Deng, Aixin Sun, and Xuying Meng. 2022b. Perplexity from plm is unreliable for evaluating text quality. *arXiv preprint arXiv:2210.05892*.
- L Winerman. 2012. Tracking the scent of information. *Monitor on psychology*, 43(3):44–47.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- xiaoju ye. 2023. [calflops: a flops and params calculate tool for neural networks in pytorch framework](#).
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate nearest neighbor negative contrastive learning for dense text retrieval. *arXiv preprint arXiv:2007.00808*.
- Wenhan Xiong, Xiang Lorraine Li, Srini Iyer, Jingfei Du, Patrick S. H. Lewis, William Yang Wang, Yashar Mehdad, Scott Yih, Sebastian Riedel, Douwe Kiela, and Barlas Oguz. 2021. [Answering complex open-domain questions with multi-hop dense retrieval](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.
- Tianjun Zhang, Shishir G Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E Gonzalez. 2024. Raft: Adapting language model to domain specific rag. *arXiv preprint arXiv:2403.10131*.

## A Datasets’ Details

In this section, we present a table that details the statistics of the datasets utilized in our study. These tables include comprehensive data such as sample sizes, feature counts, and other relevant metrics, providing a clear overview of the datasets’ composition and scope.

Table 6: Statistics of datasets.

Dataset	Train	Dev	Test
TriviaQA	78,785	8,837	11,313
NQ	79,168	8,757	3,610
WebQA	3,417	361	2,032
Entity Questions	-	-	22,000
HotpotQA	90,564	7,405	-
ArchivalQA	-	-	7,500

## B Evaluation Metrics

To assess the performance of the ASRANK, we use top-K retrieval accuracy and several other metrics for the RAG. Top-K retrieval accuracy measures whether the correct answer appears within the top-K retrieved passages, calculated as:

$$\text{TOP@}k = \frac{\sum(\text{any}(\text{Correct@}k))}{\text{Total Data}}$$

Given that LLMs tend to generate verbose answers, many standard QA metrics are not well suited to evaluate the answer quality; the Exact Match will always be less given the occurrence of other non-ground-truth tokens, and the F1 score will be penalized by other, potentially helpful tokens. Therefore, we utilize a set of model-agnostic metrics (i.e., token recall and answer string containment).

## C Retrievers

In our re-ranking experiments, we retrieve passages using both unsupervised and supervised retrievers, as detailed below.

**Unsupervised Retrievers** BM25 (Robertson et al., 2009) is a ranking function used by search engines to estimate the relevance of documents to a given search query. It is based on the probabilistic retrieval framework and uses term-frequency (TF) and inverse document frequency (IDF) of the keywords present in the question and passage.

Masked Salient Spans (MSS) (Sachan et al., 2021a) is a dense retriever trained by predicting masked salient spans like named entities with the

help of a reader network. The objective function for training the MSS retriever can be represented as:

$$\mathcal{L}_{MSS} = -\mathbb{E}_{(q,d^+,d^-)\sim D}[\log p(d^+|q) + \log(1 - p(d^-|q))]$$

where  $D$  is the dataset,  $(q, d^+, d^-)$  is a triplet of the question, positive document, and negative document, and  $p(d|q)$  is the probability of a document  $d$  being relevant to a question  $q$ .

Contriever is a framework for pre-training and fine-tuning models for information retrieval using contrastive learning. The objective function for training the Contriever model is:

$$\mathcal{L}_{Contriever} = -\mathbb{E}_{(q,d^+,d^-)\sim D}[\log \sigma(s(q, d^+)) + \log(1 - \sigma(s(q, d^-)))]$$

where  $s(q, d)$  is the similarity score between question  $q$  and document  $d$ , and  $\sigma$  is the sigmoid function (Izacard et al., 2022).

**Supervised Retrievers** Dense Passage Retrieval (DPR) (Karpukhin et al., 2020) uses annotated question-context paragraphs and hard negative examples to train a supervised dense retriever. The objective function for training the DPR model is:

$$\mathcal{L}_{DPR} = -\mathbb{E}_{(q,d^+,d^-)\sim D}[\log \sigma(s(q, d^+)) + \log(1 - \sigma(s(q, d^-)))]$$

where  $s(q, d)$  is the similarity score between question  $q$  and document  $d$ , and  $\sigma$  is the sigmoid function.

MSS-DPR (Sachan et al., 2021a) is an approach that further improves DPR performance by first pre-training the dense retriever using MSS followed by DPR-style supervised fine-tuning. The objective function for training the MSS-DPR model is:

$$\mathcal{L}_{MSS-DPR} = \alpha \mathcal{L}_{MSS} + (1 - \alpha) \mathcal{L}_{DPR}$$

where  $\alpha$  is a hyperparameter that controls the trade-off between the MSS and DPR losses.

## D Implementation Framework

Our implementation of ASRANK utilizes the PyTorch (Paszke et al., 2019) framework alongside the transformers (Wolf et al., 2019) library from Hugging Face to handle the computational demands of our document re-ranking tasks.

## E Additional Results

### E.1 The performance improvement

The performance improvement of the ASRANK is focused on using zero-shot answer scent generation with a cross-attention mechanism within its re-ranking framework. ASRANK uses the advanced capabilities of LLMs to interpret and generate answer scents. The answer scent is not static but dynamically interacts with the passage tokens through a cross-attention mechanism employed in the model’s architecture. Each token of the generated answer scent considers every token in the passage, enabling a deeper and more contextual understanding before determining the relevance of each passage. By focusing on the semantic and contextual alignment between the question and the document, ASRANK improves the retrieval and ensures that the top-ranked documents are relevant to the information needs.

### E.2 Archival Questions Result

In this section, we present the results for the ArchivalQA dataset using different retrievers, both before and after re-ranking with Llama-3-70B and GPT-3.5-turbo-0125 models combined with the T5-base model. The baseline results include performance metrics for Contriever, BM25, DPR, Rocket, and ANCE, showing their Top-1, Top-5, and Top-10 retrieval accuracy and average performance. Among the baselines, BM25 performs the best in Top-1, Top-5, Top-10 metrics, with DPR following closely.

After applying re-ranking with Llama-3-70B, all retrievers see significant improvements. The re-ranked results for DPR and ANCE particularly stand out, both achieving the highest accuracy in Top-1 (27.5% and 27.3% respectively) and Top-5/Top-10 metrics.

The results for re-ranking with GPT-3.5-turbo-0125 similarly show improvements across all metrics. ANCE continues to perform well, achieving the highest Top-1 accuracy (28.1%) after re-ranking with GPT-3.5-turbo, while DPR and BM25 also show enhanced performance compared to their baseline results.

### E.3 EntityQuestions Result

This subsection focuses on the evaluation of the Entity Questions dataset using various retrievers. The baseline models include MSS, DPR, MSS-DPR, and Contriever, and their retrieval performance is

Retriever	ArchivalQA Questions			
	Top-1	Top-5	Top-10	Avg
<i>Baselines</i>				
Contriever	1	3.2	5.0	3.0
BM25	<b>18.2</b>	<b>32.3</b>	<b>38.6</b>	<b>29.7</b>
DPR	17.0	30.1	36.8	27.9
Rocket	15.7	29.3	35.6	26.9
ANCE	18.0	31.8	37.7	29.2
<i>After Re-ranking with Llama-3-70B</i>				
Contriever+T5 <sub>base</sub>	3.9	8.1	10.4	7.4
BM25+T5 <sub>base</sub>	26.2	37.3	42.4	35.3
DPR+T5 <sub>base</sub>	<b>27.5</b>	<b>38.2</b>	<b>43.3</b>	<b>36.3</b>
Rocket+T5 <sub>base</sub>	26.2	37.4	42.4	35.3
ANCE+T5 <sub>base</sub>	27.3	<b>38.2</b>	<b>43.3</b>	<b>36.3</b>
<i>After Re-ranking with GPT-3.5-turbo-0125</i>				
Contriever+T5 <sub>base</sub>	4.2	8.7	10.9	7.9
BM25+T5 <sub>base</sub>	27.6	37.7	42.4	35.9
DPR+T5 <sub>base</sub>	27.7	38.5	43.5	36.6
Rocket+T5 <sub>base</sub>	26.5	37.9	42.7	35.7
ANCE+T5 <sub>base</sub>	<b>28.1</b>	<b>38.1</b>	<b>42.9</b>	<b>36.3</b>

Table 7: Top-1, 5, 10 retrieval accuracy for the ArchivalQA dataset, comparing baseline retrievers with results after re-ranking using Llama 70b and GPT3.5 models.

measured in terms of Top-1, Top-5, and Top-10 accuracy, as well as average accuracy across these metrics.

After applying re-ranking with Llama-3-70B, all retrievers see notable improvements in accuracy. Contriever, when combined with T5-base, achieves the highest Top-1, Top-5, and Top-10 scores, outperforming the baseline results by a large margin. The other retrievers, including MSS-DPR and DPR, also see significant gains after re-ranking, underscoring the effectiveness of using Llama-3-70B for enhancing retrieval performance.

The re-ranking with GPT-3.5-turbo-0125 produces similar improvements. Contriever again leads the performance metrics, achieving the highest Top-1 accuracy (48.9%) and consistently strong Top-5 and Top-10 results. MSS-DPR also demonstrates robust improvements, further validating the ability of the re-ranking process to enhance retrieval accuracy for the Entity Questions dataset. These results highlight the strength of combining LLM-based answer scent generation with efficient re-ranking techniques to boost retrieval performance across diverse datasets.

### E.4 Answer Scent Length

This subsection analyzes the impact of varying Answer Scent (AS) lengths on retrieval performance using the DPR on the NQ-test dataset. Table 9 presents the results of four different AS lengths: 50, 75, 100, and 128 tokens, increasing the length of the answer scent improves retrieval performance. Starting with an AS length of 50 tokens, the model

Retriever	Entity Questions			Avg
	Top-1	Top-5	Top-10	
<i>Baselines</i>				
MSS	19.3	35.9	43.1	32.8
DPR	25.3	39.5	45.3	36.7
MSS-DPR	<b>30.3</b>	47.7	54.1	<b>44.0</b>
Contriever	27.1	<b>48.0</b>	<b>55.7</b>	43.6
<i>After Re-ranking with Llama-3-70B</i>				
MSS+ $T5_{base}$	44.5	58.3	62.7	55.2
DPR+ $T5_{base}$	41.7	53.8	58.2	51.2
MSS-DPR+ $T5_{base}$	46.4	60.1	64.5	57.0
Contriever+ $T5_{base}$	<b>46.6</b>	<b>61.1</b>	<b>65.9</b>	<b>57.9</b>
<i>After Re-ranking with GPT-3.5-turbo-0125</i>				
MSS+ $T5_{base}$	46.6	60.5	64.5	57.2
DPR+ $T5_{base}$	43.6	55.6	59.4	52.9
MSS-DPR+ $T5_{base}$	48.4	62.1	66.2	58.9
Contriever+ $T5_{base}$	<b>48.9</b>	<b>63.2</b>	<b>67.5</b>	<b>59.8</b>

Table 8: Top-1, 5, 10 retrieval accuracy for the Entity Questions dataset, comparing baseline retrievers with results after re-ranking using Llama-3-70B and GPT-3.5-turbo-0125 models.

Retriever	AS length	Top-1	Top-5	Top-10	Avg
DPR	50	42.6	62.3	69.3	58.0
	75	46.5	64.5	78.2	63.0
	100	47.5	66.1	73.7	62.4
	128	51.3	70.6	76.0	65.9

Table 9: Impact of answer scent length on ranking performance using DPR on NQ-test.

achieves a Top-1 accuracy of 42.6%, Top-5 accuracy of 62.3%, and Top-10 accuracy of 69.3%, with an average accuracy of 58.0%. When the AS length increases to 75 tokens, all metrics improve, particularly the Top-10 accuracy, which jumps to 78.2%.

Further increasing the AS length to 100 and 128 tokens leads to even higher gains. With 128 tokens, the model achieves its best performance, reaching a Top-1 accuracy of 51.3%, Top-5 accuracy of 70.6%, and Top-10 accuracy of 76.0%, resulting in an overall average of 65.9%.

### E.5 Impact of Answer Scent LLM

In this section, we evaluate the impact of different LLMs on Top-1, 5, 10, 20 by re-ranking the Top-1000 passages from the NQ development set. The performance of these LLMs on the NQ development set is detailed in Table 10. The baseline retrieval using BM25 achieves a Top-1 of 22.3%. However, with LLMs like Llama-2 and GPT3.5, there’s an increase in all Top-K. For instance, Llama-2 70B improves the Top-1 accuracy to 45.3%, and GPT3.5 pushes it further to 46.3%. As the model size increases from 7B to over 70B, there’s a performance improvement. The Mixtral model achieves a Top-1 of 42.5%. The increase in Top-K from Llama-2 70B to Qwen1.5 110B is

marginal in Top-1 (from 45.3% to 44.0%). This suggests that larger models have better re-ranking capabilities due to their ability to process a broader scope of linguistic nuances. However, while larger models generally perform better, the rate of improvement varies. For example, the increase in Top-K from Llama-2 70B to Qwen1.5 110B is marginal in Top-1 (from 45.3% to 44.0%), indicating diminishing returns at higher parameter counts.

Retriever	#Parameters	NQ (dev)				
		Top-1	Top-5	Top-10	Top-20	Avg
BM25	-	22.3	43.8	53.7	62.3	
Gemma	7B	21.2	37.7	45.9	54.2	39.8
Mistral	7B	27.9	46.3	54.8	62.3	47.8
Qwen1.5	7B	30.3	50.4	58.6	66.2	51.4
Llama-2	7B	39.2	58.6	65.8	71.4	58.8
Llama3	8B	39.8	60.2	66.6	71.9	59.6
Qwen1.5	14B	34.9	54.4	62.7	69.1	55.3
Qwen1.5	32B	39.9	60.3	67.2	72.9	60.1
Mixtral	8x7B	42.5	61.9	68.2	73.0	61.4
Llama3	70B	44.8	64.1	69.9	<b>75.0</b>	<b>63.5</b>
Qwen1.5	72B	43.2	62.6	68.9	73.9	62.2
Llama-2	70B	45.3	64.0	<b>69.9</b>	74.4	63.4
Qwen1.5	110B	44.0	63.3	69.8	74.4	62.9
GPT3.5	175B	<b>46.3</b>	<b>63.6</b>	69.1	73.8	63.2

Table 10: Performance metrics of different LLMs utilizing the answer scent concept for document retrieval across Top-1, 5, 10, and 20 rankings on the NQ (dev) dataset.

### E.6 Impact of Answer Scent on Re-Ranking and Latency Implications

Answer scents improve the alignment of retrieved documents with the question. The process is captured through the computation of log-likelihood of each document given the question  $\log p(\mathbf{d} | \mathbf{q}; \Theta) = \frac{1}{|\mathbf{d}|} \sum_t \log p(d_t | \mathbf{d}_{<t})$ , where  $\mathbf{d}$  denotes the document tokens,  $\mathbf{q}$  the question,  $\Theta$  the LLM parameters, and  $|\mathbf{d}|$  the number of tokens in the document.

Re-ranking with Answer Scent has shown improvements in Top-K on the NQ development set. The Top-1 accuracy increases from 22.3% at 100 documents to around 39.8% at 750 documents (see Fig. 5). The ASRANK significantly reduces latency challenges, especially as the number of re-ranked documents grows. For example, re-ranking 1,000 documents takes up to 6.7 seconds with Llama models and 3.8 seconds with GPT models, compared to 11.6 seconds with UPR. This means ASRANK cuts latency by nearly 42% compared to UPR, as shown in Figure 5, and Table 11 in Appendix E.7.

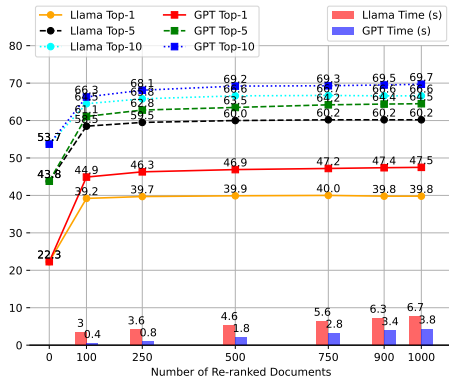


Figure 5: Effect of the number of passage candidates on the accuracy of Top-1, 5, 10 results, and latency when re-ranked with LLama 8B and GPT 175. The results were computed on the NQ development set using BM25 retrieved passages.

Retriever / Re-Ranker	#Document	NQ (dev)					Time/Question
		Top-1	Top-5	Top-10	Top-20	Top-100	
BM25	-	22.3	43.8	53.7	62.3	76.0	-
Llama3 8B	100	39.2	58.5	64.5	69.8	76.0	3s
Llama3 8B	200	39.6	59.4	65.7	70.9	78.5	3.4s
Llama3 8B	250	39.7	59.5	65.8	71.2	79.1	3.6s
Llama3 8B	500	39.9	60.0	66.6	71.9	80.2	4.6s
Llama3 8B	750	40.0	60.2	66.7	71.9	80.74	5.6s
Llama3 8B	900	39.8	60.2	66.6	72.0	80.9	6.3s
Llama3 8B	1000	39.8	60.2	66.6	71.9	80.9	6.7
GPT 175B	100	44.9	61.1	66.3	70.8	76.0	0.4s
GPT 175B	250	46.3	62.8	68.1	72.4	79.4	0.8s
GPT 175B	500	46.9	63.5	69.2	73.6	80.8	1.8s
GPT 175B	750	47.2	64.2	69.3	74.2	81.6	2.8s
GPT 175B	900	47.4	64.4	69.5	74.3	81.8	3.4s
GPT 175B	1000	47.5	64.5	69.7	74.3	81.9	3.8s

Table 11: Impact of the Number of Passage Candidates on Top-1, Top-5, Top-10 Retrieval Accuracy, and Latency per Query.

### E.7 Impact of Passage Number on Retrieval Accuracy and Latency

In this section, we analyze the relationship between the number of passages re-ranked and both retrieval accuracy and latency. This study highlights how the ASRANK performs as we increase the number of passage candidates, focusing on Top-K retrieval accuracy and the time taken per query. We conducted experiments using the NQ development set to evaluate the performance of ASRANK with different quantities of retrieved passages. The passages were retrieved using BM25 and re-ranked using LLama (8B) and GPT (175B) models. We varied the number of passages from 100 to 1000 to observe the impact on Top-K accuracy and latency. The results of these experiments are presented in Table 11. The table illustrates how increasing the number of re-ranked passages affects the Top-1, 5, 10 retrieval metrics, and the latency per query.

**Retrieval Accuracy:** The Top-1 accuracy significantly improves as the number of re-ranked passages increases. For example, using Llama3 8B, Top-1 accuracy increases from 39.2% with 100 passages to 40.0% with 750 passages. Similarly, GPT 175B shows an increase in Top-1 accuracy from 44.9% with 100 passages to 47.5% with 1000 passages.

**Latency:** As expected, the latency per query increases with the number of passages. With Llama3 8B, the latency grows from 3 seconds for 100 passages to 6.7 seconds for 1000 passages. GPT 175B, while providing better accuracy, also shows an increase in latency, from 0.4 seconds for 100 passages to 3.8 seconds for 1000 passages.

### E.8 Comparative Analysis of LLama 7B and UPR for Document Re-Ranking

In this section, we present a comparison between the performance of ASRANK utilizing the LLama 7B model and the UPR method. This analysis is aimed at understanding how ASRANK, enhanced with the capabilities of LLama 7B, measures up against UPR in terms of improving retrieval accuracy across various question-answering datasets.

We evaluated both LLama 7B with ASRANK and UPR across three major datasets: NQ, TriviaQA, and WebQ. The goal was to assess the improvements in retrieval accuracy, specifically focusing on Top-1, Top-5, and Top-10 metrics. The retrieval setups included unsupervised and supervised retrievers. The detailed results are summarized in Table 12. The analysis highlights the performance of the two methods under different retrievers, providing insights into their effectiveness across varying retrieval conditions.

**Performance across Datasets:** Both methods improve retrieval accuracy across all datasets. However, ASRANK with LLama 7B consistently achieves a higher Top-1 metric compared to UPR, suggesting that the inclusion of the answer scent concept might be more effective at distinguishing the most relevant documents at the top of the retrieval list.

**Influence of Retrieval Method:** When combined with MSS, ASRANK with LLama 7B surpasses UPR in Top-1 retrieval accuracy by a notable margin (e.g., 41.3% vs. 38.7% on NQ). This indicates that ASRANK’s approach to utilizing deep contextual embeddings effectively captures

Retriever	NQ				TriviaQA				WebQ			
	Top-1	Top-5	Top-10	Avg	Top-1	Top-5	Top-10	Avg	Top-1	Top-5	Top-10	Avg
<i>Unsupervised Retrievers</i>												
MSS	19.2	41.2	51.2	37.2	30.7	52.6	60.5	47.9	11.6	29.0	39.1	26.6
MSS + UPR	38.7	<b>64.8</b>	<b>72.2</b>	<b>58.6</b>	57.2	<b>75.5</b>	78.9	70.5	29.9	57.4	65.0	50.7
MSS + ASRANK + $T5_{base}$	41.3	60.3	67.2	56.2	58.5	71.8	75.6	68.6	40.1	59.9	66.6	55.5
MSS + ASRANK + $T5_{Large}$	<b>42.8</b>	63.2	70.0	<b>58.6</b>	<b>59.1</b>	<b>76.4</b>	<b>79.2</b>	<b>71.5</b>	<b>40.4</b>	<b>60.2</b>	<b>66.9</b>	<b>55.8</b>
BM25	22.1	43.7	54.4	40.1	46.3	66.2	71.7	61.4	18.8	41.8	52.1	37.6
BM25 + UPR	35.4	63.4	70.2	56.3	55.7	76.5	<b>80.2</b>	70.8	30.1	57.3	66.5	51.3
BM25 + ASRANK + $T5_{base}$	42.1	61.1	67.4	56.8	58.2	71.1	74.7	68.0	40.9	61.1	68.1	56.7
BM25 + ASRANK + $T5_{Large}$	<b>44.3</b>	<b>64.2</b>	<b>71.0</b>	<b>59.8</b>	<b>60.8</b>	<b>77.0</b>	80.1	<b>72.6</b>	<b>41.2</b>	<b>62.3</b>	<b>68.2</b>	<b>57.2</b>
Contriever	22.1	47.2	58.7	42.7	34.1	59.4	68.0	53.8	19.9	43.4	56.3	39.9
Contriever + UPR	36.4	<b>64.6</b>	<b>72.4</b>	<b>57.8</b>	56.7	<b>77.0</b>	<b>80.2</b>	<b>71.3</b>	30.0	58.5	68.2	52.2
Contriever + ASRANK + $T5_{base}$	<b>41.5</b>	61.3	68.4	57.0	<b>57.9</b>	72.8	76.8	69.1	<b>42.9</b>	<b>62.7</b>	<b>69.8</b>	<b>58.4</b>
<i>Supervised Retrievers</i>												
DPR	48.6	68.7	74.5	63.9	57.4	72.4	76.5	68.7	44.8	65.0	70.6	60.1
DPR + UPR	42.5	<b>70.6</b>	<b>78.1</b>	<b>63.8</b>	61.3	<b>78.7</b>	<b>81.9</b>	<b>74.0</b>	34.9	63.6	71.7	56.7
DPR + ASRANK + $T5_{base}$	<b>43.5</b>	64.9	72.2	60.2	<b>61.8</b>	74.6	78.3	71.5	<b>45.9</b>	<b>66.7</b>	<b>72.4</b>	<b>61.6</b>
MSS-DPR	<b>50.1</b>	<b>71.8</b>	<b>77.4</b>	<b>66.5</b>	61.6	75.2	79.1	71.9	44.2	65.0	71.6	60.3
MSS-DPR + UPR	41.4	69.8	77.9	63.0	60.5	<b>78.9</b>	<b>82.5</b>	<b>74.0</b>	31.8	61.6	70.3	54.5
MSS-DPR + ASRANK + $T5_{base}$	43.5	65.1	72.5	60.3	<b>61.7</b>	74.8	78.6	71.7	<b>44.6</b>	<b>65.4</b>	<b>72.2</b>	<b>60.7</b>

Table 12: Top-1, 5, 10 retrieval accuracy on the test set of datasets before and after re-ranking the top 1000 retrieved passages. § refers to Llama 7B

nuances that improve the alignment between the query and retrieved documents.

## E.9 FLOPs and Latency Comparison

To evaluate the efficiency of our approach, we compute the total FLOPs and latency required for reranking 1,000 passages, comparing our method with the UPR method. The FLOPs were calculated using the fvcore and calcflops library (xiaoju ye, 2023).

**UPR Method:** The total number of FLOPs for reranking 1,000 passages is approximately  $2 \times 10^{15}$  FLOPs.

**ASRANK:** The total number of FLOPs required is significantly lower, at approximately  $1.1 \times 10^{12} + 1.1 \times 10^{15} \approx 1.1011 \times 10^{15}$  FLOPs. This is because our approach generates the answer scent using a larger model (Llama 8B), but only once for the question. Subsequently, the reranking process is handled by a smaller model, such as T5 base, which uses the generated answer scent to rerank the passages.

For latency comparisons, we measured the time needed to rerank 1,000 passages based on the computational times of GPT, Llama 3-8B, and UPR T0 3B models. For our method, Llama 3-8B is used once to generate the answer scent in 2.77 seconds, and T5 base requires 2.54 seconds to rerank 1,000 passages, leading to a total of 5.31 seconds. In contrast, UPR T0 3B takes an average of 6.7 seconds to rerank the same number of passages. These results highlight the efficiency of our approach, as it reduces both the computational cost (FLOPs) and the reranking latency compared to UPR.

## F RAG

In the realm of Retrieval-Augmented Generation (RAG), our study delves into the effects of utilizing LLaMA 7B and LLaMA 13B models, along with varying the number of documents considered in the re-ranking process. Our examination reveals differences in performance across two scenarios: using either one or two documents during the re-ranking phase.

Starting with the LLaMA 7B model, we observed that increasing the number of documents from one to two generally improves the recall and contextual understanding of the model, which is critical in generating accurate responses. For instance, when using the MSS-DPR retriever with LLaMA 7B, the exact match (EM) score sees a slight improvement from 24.3% with one document to 24.9% with two documents. This pattern is consistent across other retrievers like BM25 and Contriever, suggesting that the additional context from a second document helps the model refine its answers.

Switching to the LLaMA 13B model, which offers more capacity and potentially finer understanding due to its larger size. For example, when using the BM25 retriever with LLaMA 13B, the EM score increases from 18.5% to 28.8% with two documents. This suggests that the larger model can leverage the extra information more effectively, leading to better overall performance.



## G Case Study

In this section, we present a detailed case study to illustrate the effectiveness of ASRANK in re-ranking documents retrieved by different retrieval systems. Tables 14, 15, and 16 showcase examples from the NQ dev dataset, WebQA, and TriviaQA, respectively. Each table lists the document IDs retrieved before and after applying ASRANK, indicating whether each document contains the answer ('has\_answer: True' or 'False'). These case studies demonstrate how ASRANK enhances the precision of document retrieval across varied contexts and query types by leveraging the answer scent generated from advanced language models.

Model	top- <i>K</i>	NQ			TriviaQA			WebQA		
		EM	Recall	Con	EM	Recall	Con	EM	Recall	Con
<b>LLama-2 7B</b>										
Question Only	0	14.4	28.6	21.6	41.7	54.6	48.6	14.8	40.4	30.3
<b>LLama-2 13B</b>										
Question Only	0	11.3	27.7	20.6	39.1	55.2	48.8	11.9	42.8	30.2
<b>LLama-2 7B+Baselines</b>										
BM25	1	16.0	29.3	21.7	51.9	63.5	57.2	14.3	35.7	25.6
MSS	1	14.9	27.4	20.8	43.6	55.4	49.3	13.7	37.1	26.9
Contriever	1	17.3	31.1	23.9	44.4	56.5	50.2	13.6	38.5	23.8
DPR	1	24.6	<b>40.5</b>	<b>32.1</b>	50.0	62.6	56.6	15.1	<b>40.3</b>	<b>29.3</b>
MSS-DPR	1	<b>24.9</b>	40.4	32.0	<b>51.9</b>	<b>64.7</b>	<b>58.4</b>	<b>15.9</b>	40.1	29.0
<b>LLama-2 13B+Baselines</b>										
BM25	1	18.5	30.8	23.9	54.6	65.8	59.4	14.6	35.1	25.5
MSS	1	17.9	29.5	22.7	47.8	58.5	52.6	14.9	35.6	25.5
Contriever	1	20.3	32.4	25.2	49.0	59.5	53.6	17.5	38.6	28.3
DPR	1	27.9	43.2	33.9	53.6	65.3	58.9	18.2	<b>41.2</b>	<b>30.9</b>
MSS-DPR	1	<b>28.9</b>	<b>43.8</b>	<b>34.5</b>	<b>54.7</b>	<b>66.7</b>	<b>60.5</b>	<b>19.6</b>	40.7	30.7
<b>LLama-1 7B+Baselines</b>										
MSS	2	15.1	28.2	21.2	44.6	57.3	51.3	14.3	37.8	27.2
BM25	2	16.3	29.9	22.8	52.2	65.0	58.5	13.1	36.9	26.1
Contriever	2	16.9	31.3	23.9	44.8	58.2	51.9	13.9	39.3	28.4
DPR	2	23.9	39.3	31.5	49.8	63.3	57.0	14.8	<b>40.3</b>	<b>28.7</b>
MSS-DPR	2	<b>24.3</b>	<b>40.3</b>	<b>32.2</b>	<b>50.8</b>	<b>64.5</b>	<b>58.1</b>	<b>15.1</b>	39.7	29.0
<b>LLama-2 7B+UPR</b>										
MSS	1	21.9	37.3	29.2	53.5	66.5	60.1	15.1	39.9	<b>29.0</b>
BM25	1	21.2	36.2	28.4	<b>57.6</b>	<b>70.3</b>	<b>63.6</b>	13.9	37.0	25.6
Contriever	1	22.5	38.5	30.5	53.8	67.5	61.3	14.0	38.9	27.6
DPR	1	<b>23.8</b>	<b>39.8</b>	<b>31.3</b>	55.2	68.5	61.8	<b>15.5</b>	<b>40.1</b>	28.7
MSS-DPR	1	23.6	39.4	30.8	55.3	68.5	62.2	15.3	39.8	28.4
<b>LLama-2 13B+UPR</b>										
MSS	1	25.2	39.4	31.2	56.4	68.4	62.1	16.9	39.3	28.0
BM25	1	25.1	39.3	30.7	57.3	68.8	<b>63.5</b>	16.8	36.8	26.7
Contriever	1	26.0	40.4	31.9	56.5	68.0	62.7	17.4	38.3	28.4
DPR	1	<b>27.4</b>	<b>42.2</b>	<b>33.0</b>	<b>57.3</b>	<b>69.6</b>	63.2	<b>17.5</b>	<b>40.6</b>	<b>29.8</b>
MSS-DPR	1	26.3	41.3	32.7	57.2	69.2	62.9	17.1	37.9	27.2
<b>LLama-1 7B+UPR</b>										
MSS	2	21.6	37.3	29.9	54.1	67.9	61.3	15.2	39.1	28.1
BM25	2	22.0	37.8	30.2	<b>58.2</b>	<b>71.4</b>	<b>64.8</b>	14.7	39.7	28.1
Contriever	2	22.3	38.4	30.5	54.9	68.2	65.0	14.5	38.8	27.2
DPR	2	23.2	38.9	31.3	55.1	69.3	62.9	<b>15.7</b>	<b>40.4</b>	<b>28.8</b>
MSS-DPR	2	<b>24.1</b>	<b>40.0</b>	<b>32.0</b>	54.9	69.2	62.7	14.4	39.6	27.5
<b>LLama-2 7B+ASRANK</b>										
MSS	1	24.8	40.6	32.7	57.1	70.5	64.1	17.9	42.3	31.2
BM25	1	25.0	40.4	32.3	<b>60.6</b>	<b>73.2</b>	<b>66.8</b>	16.9	42.5	31.6
Contriever	1	25.9	41.9	33.6	57.8	71.0	64.9	17.7	43.9	33.1
DPR	1	25.8	42.2	33.7	57.6	71.1	64.6	16.5	43.7	31.2
MSS-DPR	1	<b>25.9</b>	<b>42.6</b>	<b>34.2</b>	58.9	71.8	65.4	<b>18.1</b>	<b>43.9</b>	<b>32.8</b>
<b>LLama-2 13B+ASRANK</b>										
MSS	1	28.5	43.4	34.6	60.1	72.4	65.9	20.5	43.9	33.6
BM25	1	28.8	44.2	35.4	<b>63.3</b>	<b>74.9</b>	<b>68.5</b>	19.3	43.0	31.5
Contriever	1	29.7	45.1	36.3	60.1	72.4	66.1	20.6	44.0	33.2
DPR	1	28.9	44.9	35.5	60.8	72.9	66.7	19.9	43.0	32.2
MSS-DPR	1	<b>30.2</b>	<b>45.5</b>	<b>36.6</b>	60.9	73.3	66.9	<b>20.6</b>	<b>44.7</b>	<b>33.7</b>
<b>LLama-1 7B+ASRANK</b>										
MSS	2	25.5	41.3	33.2	56.5	70.8	64.2	17.9	43.8	32.9
BM25	2	24.9	40.8	33.4	<b>59.8</b>	<b>73.3</b>	<b>67.1</b>	<b>16.6</b>	<b>41.7</b>	<b>30.4</b>
Contriever	2	25.5	41.6	33.8	56.5	71.4	64.9	17.0	43.4	31.8
DPR	2	25.9	41.9	33.8	57.3	71.2	64.7	17.6	43.6	32.3
MSS-DPR	2	<b>26.0</b>	<b>42.3</b>	<b>34.4</b>	57.7	72.1	65.5	17.6	43.7	32.5

Table 13: Exact match scores for the open-domain QA task.

---

**Answer Scent Prompt:** Generate a brief, insightful answer scent to the following question: *q*

---

**Question:** who sang i just called to say i love you?

**Answer Scent:** The song "I Just Called to Say I Love You" was sung by Stevie Wonder. It was released in 1984 and was featured in the romantic comedy film "The Woman in Red". The song was a huge commercial success, reaching number one on the Billboard Hot 100 chart and winning several awards, including the Academy Award for Best Original Song.

**True Answer:** Stevie Wonder

**DPR:**

6069735, has\_answer: True  
6069737, has\_answer: True  
8421731, has\_answer: False  
8421729, has\_answer: False  
8421736, has\_answer: False

**DPR+ASRANK:**

6069735, has\_answer: True  
6069737, has\_answer: True  
12952376, has\_answer: True  
5693954, has\_answer: False  
6069736, has\_answer: False

---

**Question:** who is the coach of arizona state men's basketball?

**Answer Scent:** the head coach of the Arizona State University men's basketball team is Bobby Hurley.

**True Answer:** Bobby Hurley

**DPR:**

12295065, has\_answer: True  
20172309, has\_answer: False  
10047172, has\_answer: False  
10047129, has\_answer: False  
17545880, has\_answer: False

**DPR+ASRANK:**

4833080, has\_answer: True  
4833089, has\_answer: True  
20047916, has\_answer: True  
20776179, has\_answer: True  
20776180, has\_answer: True

---

**Question:** who played mason on wizards of waverly place?

**Answer Scent:** The character Mason Greyback, a werewolf and love interest of Alex Russo, was played by actor Gregg Sulkin on the Disney Channel show "Wizards of Waverly Place".

**True Answer:** Sulkin

**DPR:**

4309105, has\_answer: False  
15243289, has\_answer: False  
19878866, has\_answer: False  
17449330, has\_answer: False  
5763564, has\_answer: False

**DPR+ASRANK:**

10341922, has\_answer: True  
9651553, has\_answer: True  
17449330, has\_answer: False  
17449350, has\_answer: False  
17449363, has\_answer: False

---

**Question:** who plays the queen on game of thrones?

**Answer Scent:** The role of Cersei Lannister, the Queen Regent of the Seven Kingdoms, on the hit HBO series Game of Thrones is played by the talented actress Lena Headey.

**True Answer:** Headey

**DPR:**

15742943, has\_answer: False  
16420109, has\_answer: False  
13408263, has\_answer: False  
15179481, has\_answer: False  
13408259, has\_answer: False

**DPR+ASRANK:**

4456301, has\_answer: True  
5396490, has\_answer: True  
4456267, has\_answer: True  
5396478, has\_answer: True  
12800938, has\_answer: True

---

Table 14: Case study of the retrieved passage from DPR retriever before/after reranking for NQ dev dataset.

---

**Answer Scent Prompt:** Generate a brief, insightful answer scent to the following question: *q*

---

**Question:** where is jamarcus russell from?

**Answer Scent:** JaMarcus Russell, the former NFL quarterback, is from Mobile, Alabama.

**True Answer:** Mobile

**MSS:**

7564117, has\_answer: False  
9565056, has\_answer: False  
7564095, has\_answer: False  
7564124, has\_answer: True  
8298815, has\_answer: False

**MSS+ASRANK:**

7564126, has\_answer: True  
7564096, has\_answer: True  
7564097, has\_answer: True  
7564125, has\_answer: True  
7564127, has\_answer: False

---

**Question:** where is the fukushima daiichi nuclear plant located?

**Answer Scent:** The Fukushima Daiichi Nuclear Power Plant is located in Okuma, Fukushima Prefecture, Japan. It is situated on the Pacific coast of Honshu, the largest island of Japan, approximately 150 miles (240 km) north of Tokyo.

**True Answer:** "Japan", "Okum"

**MSS:**

20927197, has\_answer: False  
3072697, has\_answer: False  
20927199, has\_answer: False  
10806688, has\_answer: False  
18499449, has\_answer: True

**MSS+ASRANK:**

15436005, has\_answer: False  
11148375, has\_answer: True  
10483298, has\_answer: True  
3541428, has\_answer: False  
1682872, has\_answer: True

---

**Question:** what does jamaican people speak?

**Answer Scent:** In Jamaica, the official language is English, which is used in government, education, business, and formal settings. However, the most widely spoken language in Jamaica is Jamaican Patois, also known as Jamaican Creole or Patwa.

**True Answer:** Jamaican English

**MSS:**

5665719, has\_answer: False  
9912963, has\_answer: False  
5665720, has\_answer: False  
11838832, has\_answer: False  
20587290, has\_answer: False

**MSS+ASRANK:**

4423284, has\_answer: False  
1353789, has\_answer: False  
8404038, has\_answer: False  
4423299, has\_answer: False  
4423301, has\_answer: False

---

**Question:** what is the best sandals resort in st lucia?

**Answer Scent:** St. Lucia is a beautiful island with several amazing Sandals Resorts to choose from. Each resort has its unique features, amenities, and atmosphere, so the "best" one ultimately depends on your personal preferences and priorities.

**True Answer:** "Micoud Quarter", "Choiseul Quarter", "Praslin Quarter", ..

**MSS:**

18392196, has\_answer: False  
18461202, has\_answer: False  
11371584, has\_answer: False  
16577459, has\_answer: False  
3764126, has\_answer: False

**MSS+ASRANK:**

5476353, has\_answer: False  
18392196, has\_answer: False  
3401309, has\_answer: False  
3401311, has\_answer: True  
6134966, has\_answer: False

---

Table 15: Case study of the retrieved passage from MSS retriever Before/after Reranking for WebQA.

---

**Answer Scent Prompt:** Generate a brief, insightful answer scent to the following question: *q*

---

**Question:** which 70s show was based on the british show till death us do part?

**Answer Scent:** The 1970s show based on the British show "Till Death Us Do Part" is "All in the Family".

**True Answer:** "All In The Family", "Justice For All (TV pilot)", "Stretch Cunningham", ...

**Contriever:**

9539720, has\_answer: False

6899634, has\_answer: False

475319, has\_answer: False

9549805, has\_answer: False

475315, has\_answer: False

**Contriever+ASRANK:**

9607452, has\_answer: True

1413988, has\_answer: True

1834891, has\_answer: True

5285410, has\_answer: True

1941863, has\_answer: True

---

**Question:** what is the name of terence and shirley conran's dress designer son?

**Answer Scent:** Jasper Conran!

**True Answer:** "Jaspis", "Bruneau jasper", "Egyptian jasper"

**Contriever:**

4935862, has\_answer: False

4935861, has\_answer: True

7176709, has\_answer: False

14139592, has\_answer: False

5848573, has\_answer: True

**Contriever+ASRANK:**

5848571, has\_answer: True

5848575, has\_answer: False

5848577, has\_answer: True

5848576, has\_answer: False

5848573, has\_answer: True

---

**Question:** in which country is the sky train rail bridge?

**Answer Scent:** The SkyTrain Rail Bridge is located in Vancouver, British Columbia, Canada.

**True Answer:** "Canada", "Kenadian", "Canadialand", "Xanada", "Dominion of Canada", "Canadaa"

**Contriever:**

11617523, has\_answer: False

11617522, has\_answer: False

7697355, has\_answer: False

3375880, has\_answer: False

4904611, has\_answer: True

**Contriever+ASRANK:**

8509738, has\_answer: True

1145807, has\_answer: True

1145854, has\_answer: True

1145799, has\_answer: True

8509740, has\_answer: True

---

**Question:** bandar seri begawan international airport is in which country?

**Answer Scent:** Bandar Seri Begawan International Airport is located in Brunei.

**True Answer:** "Abode of Peace", "BRUNEI", "Health in Brunei", ...

**Contriever:**

2693267, has\_answer: False

6595413, has\_answer: False

10932719, has\_answer: False

670520, has\_answer: True

10932726, has\_answer: True

**Contriever+ASRANK:**

670503, has\_answer: True

670496, has\_answer: True

10893158, has\_answer: True

5225731, has\_answer: True

11964123, has\_answer: True

---

Table 16: Case study of the retrieved passage from Contriever retriever Before/after Reranking for TriviaQA.