

# Inferential Question Answering

Jamshid Mozafari\*  
University of Innsbruck  
Innsbruck, Austria  
jamshid.mozafari@uibk.ac.at

Hamed Zamani  
University of  
Massachusetts Amherst  
Amherst, United States  
zamani@cs.umass.edu

Guido Zuccon  
The University of  
Queensland  
Brisbane, Australia  
g.zuccon@uq.edu.au

Adam Jatowt  
University of Innsbruck  
Innsbruck, Austria  
adam.jatowt@uibk.ac.at

## Abstract

Despite extensive research on a wide range of question answering (QA) systems, most existing work focuses on answer containment—i.e., assuming that answers can be directly extracted and/or generated from documents in the corpus. However, some questions require *inference*, i.e., deriving answers that are not explicitly stated but can be inferred from the available information. We introduce **Inferential QA**—a new task that challenges models to infer answers from *answer-supporting* passages which provide only clues. To study this problem, we construct **QUIT** (Q**U**estions requiring **I**nference from **T**exts) dataset, comprising 7,401 questions and 2.4M passages built from high-convergence human- and machine-authored hints, labeled across three relevance levels using LLM-based answerability and human verification. Through comprehensive evaluation of retrievers, rerankers, and LLM-based readers, we show that methods effective on traditional QA tasks struggle in inferential QA: retrievers underperform, rerankers offer limited gains, and fine-tuning provides inconsistent improvements. Even reasoning-oriented LLMs fail to outperform smaller general-purpose models. These findings reveal that current QA pipelines are not yet ready for inference-based reasoning. Inferential QA thus establishes a new class of QA tasks that move towards understanding and reasoning from indirect textual evidence.

## CCS Concepts

• Information systems → Retrieval models and ranking; Retrieval tasks and goals; Evaluation of retrieval results.

## Keywords

Inferential Question Answering, Hints, Retrieval-Augmented Reasoning, Reasoning, Retrieval-Augmented Generation, Large Language Models, Evaluation

### ACM Reference Format:

Jamshid Mozafari, Hamed Zamani, Guido Zuccon, and Adam Jatowt. 2026. Inferential Question Answering. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3774904.3792653>

\*Corresponding Author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WWW '26, Dubai, United Arab Emirates*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2307-0/2026/04  
<https://doi.org/10.1145/3774904.3792653>

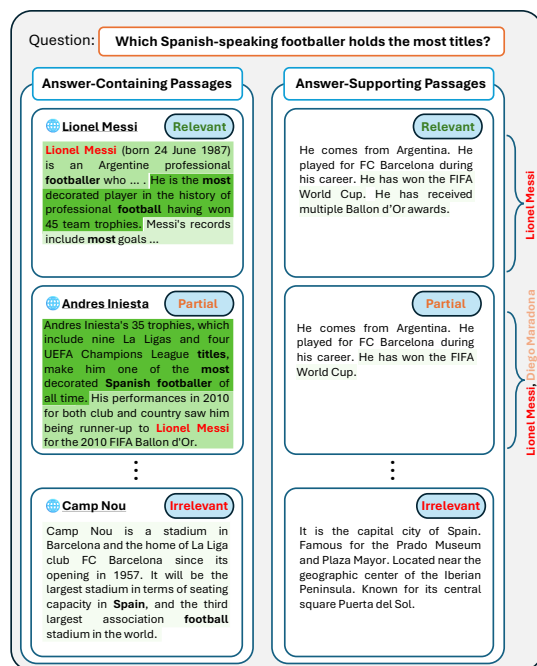


Figure 1: Example of answer-containing passages and answer-supporting passages for the question *Which Spanish-speaking footballer holds the most titles?*, whose answer is **Lionel Messi**. Green highlights indicate semantic similarity between the question and sentences (darker green = higher similarity). Bold words mark lexical overlaps between the question and passages. Blue rounded rectangles denote the relevance level of each passage to the question. In the *Answer Supporting* column, the entities that each passage implicitly describes with respect to the question are shown on the right side of the passages.

## 1 Introduction

Question Answering (QA) systems [2] aim to provide direct responses to natural language questions. Existing QA tasks include but are not limited to factoid QA [54], descriptive (long-form) and non-factoid QA [18, 27], Boolean QA [60], and multi-hop QA [59]. QA systems have been studied for decades [38] and regardless of which of these tasks they focus on, earlier QA systems primarily focused on *extractive* approaches, which locate an answer span directly in the retrieved passage. More recent QA systems emphasize on *generative* approaches, where the answer is generated conditioned on the passage content, mostly using large language

models (LLMs). Extractive QA systems relied on encoder-based models such as BERT [11] and RoBERTa [30]. However, with the advent of decoder-based generative models—including LLMs, such as LLaMA [13] and Gemma [53]—the paradigm has shifted toward Retrieval-Augmented Generation (RAG) [28]. This transition is driven by the superior generation and reasoning capabilities of LLMs, which enable them to outperform extractive systems on a broad range of natural language tasks [32].

Despite the impressive performance of LLMs on many existing QA benchmarks, we believe that there exists a relatively understudied category of QA tasks that requires further attention in the context of RAG and reasoning—the task of **Inferential Question Answering (Inferential QA)**. In this task, the answers to the questions do not exist in the corpus; instead, the model must *infer* the correct answer from contextual clues, background knowledge, and logical reasoning supported by the document. For instance, as illustrated in Figure 1, given the question “Which Spanish-speaking footballer holds the most titles?”, the answer-containing relevant passage explicitly states that *Lionel Messi* is the most decorated player, making it directly answerable. In contrast, the answer-supporting relevant passage provides several descriptive clues—that the player is from Argentina, played for FC Barcelona, won the FIFA World Cup, and received multiple Ballon d’Or awards—from which the model must *infer* that the answer is *Lionel Messi*. This task has real-world applications, such as knowledge-based reasoning systems [55] and educational tutoring [4], where deriving answers from indirect or incomplete evidence is essential for evaluating comprehension and reasoning. Furthermore, Inferential QA assesses the reasoning abilities of LLMs and RAG models in settings where answers are not stated explicitly but can be inferred from contextual evidence. Crucially, it differs from commonsense reasoning QA [50], which relies on generic world knowledge. It also differs from multi-hop QA, as it does not require putting multiple passages and entities together to extract or generate the answer.

We first construct **QUIT (QUESTions requiring Inference from Texts)**,<sup>1</sup> a dataset with dedicated training, development, and test splits, containing 7,401 questions and 2,405,325 passages. The passages in QUIT are derived from **hints** [20, 21], i.e., clues that guide individuals toward a solution without explicitly providing the answer. By concatenating multiple hints related to an entity, we create passages that describe the entity indirectly. For example, in Figure 1, the relevant, partial, and irrelevant passages in the Answer-Supported column are formed by concatenating four hints about *Lionel Messi*, three hints about *Lionel Messi* and *Diego Maradona*, and four hints about *Madrid*, respectively. Prior work [33] demonstrated that sentence-level hints improve LLM performance in QA compared to retrieved or generated passages, motivating their use in constructing our benchmark. Each passage is annotated with one of three relevance labels: **2 (fully relevant)** if the passage enables LLM to generate the correct answer, **1 (partially relevant)** if it references the answer but lacks sufficient information for correct generation, and **0 (irrelevant)** if it does not support the answer. Annotation details are presented in Section 3.2.2.

Following the QA literature that uses a Retriever, Reranker, and Reader pipeline for answer extraction or generation, we evaluate

diverse retrievers, rerankers, and readers on the Inferential QA task. Our experiments demonstrate that methods effective on existing QA datasets struggle under this setting. In particular, current retrievers and rerankers perform poorly on the Inferential QA task, and fine-tuning yields only limited or inconsistent improvements. These findings highlight that existing QA pipelines are insufficient for inferential questions and retrieval of answer-supporting passages.

To summarize, our main contributions are as follows:

- We introduce **Inferential QA**, a new and challenging QA task that requires models to infer answers from textual clues rather than extract them verbatim, highlighting a gap in current QA paradigms and the limitations of existing pipelines.
- We construct **QUIT (QUESTions requiring Inference from Texts)**, a large-scale benchmark consisting of 7,401 questions and 2,405,325 passages<sup>2</sup> derived from hints. The dataset provides dedicated training, development, and test splits, with passages labeled across three levels of relevance.
- We comprehensively evaluate retrievers, rerankers, and LLM readers on Inferential QA. Our findings show that (i) methods effective on traditional QA datasets fail to generalize to inference-based reasoning, and (ii) fine-tuning existing models yields only limited or inconsistent improvements.

We believe **Inferential QA** opens a promising direction for retrieval-augmented reasoning models, e.g., Search-R1 [22], ReasonIR [47], RaDeR [10], and DIVER [31], that are capable of *inferring* answers from evidence rather than extracting them, moving QA closer to genuine reasoning and comprehension.

## 2 Related Work

### 2.1 QA Datasets

QA research has been driven by large-scale datasets. Early benchmarks such as SQuAD [42], TriviaQA [23], WebQuestions [5], Natural Questions (NQ) [26], and BoolQ [9] focus on *answer containment*, where the passage explicitly states the answer. Reasoning-oriented datasets later emerged, including HotpotQA [59] for multi-hop reasoning, DROP [12] for numerical reasoning, and StrategyQA [16] for implicit reasoning. More recent efforts, such as ReasonChainQA [61], encourage complex inference, while BRIGHT [48] emphasizes reasoning-aware retrieval. However, these datasets still assume the presence of passages that explicitly contain the answer. In contrast, QUIT, focuses on answer-supporting passages, where only indirect evidence is available and answers must be inferred.

### 2.2 RAG

The emergence of LLMs has shifted QA toward generative paradigms, most notably RAG [28]. To improve reasoning in this setting, prior work has explored multi-hop retrieval [41, 59], iterative retrieval generation cycles [57], and prompting strategies that elicit step-by-step reasoning [56]. Despite these advances, RAG methods generally assume that retrieved passages explicitly contain the answer. By contrast, QUIT evaluates QA systems in scenarios where passages provide only indirect evidence, requiring models to infer answers from contextual clues.

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

<sup>2</sup>Throughout this paper, unless otherwise specified, passage refers to Answer-Supporting passages; Answer-Containing passages are always named explicitly.

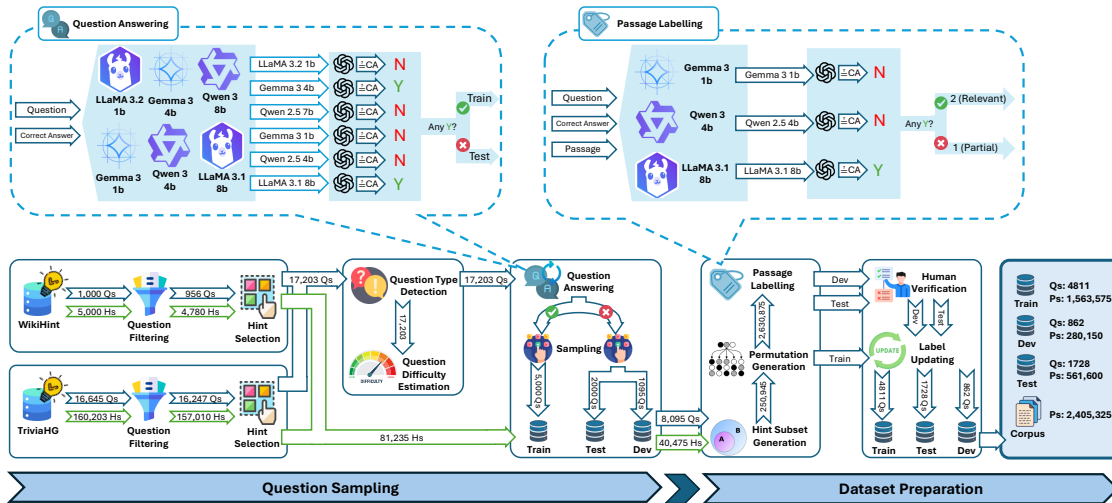


Figure 2: The pipeline for constructing the QUIT dataset. (1) **Question Sampling:** Questions (Qs) and their hints (Hs) are first filtered to avoid answer leakage. For each question, the top-5 hints are selected based on convergence, followed by question type detection and difficulty estimation. Valid questions are then sampled into temporary training, development, and test sets. (2) **Dataset Preparation:** For each question, subsets and permutations of hints are generated to form diverse passages (Ps). These passages are labeled automatically, with dev and test labels further verified and updated through human verification. The final corpus includes 7,401 questions and 2,405,325 passages across training, development, and test splits.

### 2.3 Inferential Questions

Hints have been recently studied as a way to guide LLMs without directly revealing answers [33, 36]. Such evidence connects to the notion of *inferential questions*, whose answers cannot be extracted verbatim from given texts but must be inferred from indirect clues, background knowledge, or reasoning. Beyond NLP, inferential questioning has long been emphasized in linguistics and education as central to comprehension and higher-order reasoning [17]. Within NLP, related work has addressed reasoning via multi-hop datasets [59] or implicit commonsense [16], yet these still assume that supporting passages explicitly contain the answer. To our knowledge, no dataset benchmarks QA pipelines on passages that only provide indirect evidence. QUIT addresses this gap by constructing large-scale passages from concatenated hints, enabling study of Inferential QA with retrievers, rerankers, and readers.

## 3 Methodology

This section describes the QUIT construction pipeline, including *Question Sampling* and *Dataset Preparation*, as shown in Figure 2.

### 3.1 Question Sampling

In this stage, we identify and filter valid questions and their corresponding hints to construct the foundation of the QUIT benchmark.

**3.1.1 Datasets.** To construct the QUIT benchmark, we use two existing resources that provide question–hint pairs: TriviaHG [35] and WikiHint [34]. TriviaHG contains 16,645 questions and 160,203 automatically generated hints, created using Microsoft Copilot<sup>3</sup>, with

<sup>3</sup><https://copilot.microsoft.com/>

each question associated with approximately 10 hints. In contrast, WikiHint consists of 1,000 questions and 5,000 human-written hints. A key feature shared by both datasets is that their hints are designed to describe a concept without explicitly naming it—precisely the type of indirect information needed to construct *answer-supporting passages*. These hints serve as the foundation for generating passages that imply, rather than state, the correct answer.

**3.1.2 Question Filtering.** We begin by filtering out questions for which at least one hint leaks the answer. To identify such cases, we use BEM [6] to determine whether any words in the associated hints are semantically equivalent to the correct answer. If at least one word is deemed equivalent, the question is discarded to prevent answer leakage. Next, for each remaining question, we rank its hints based on their *convergence* [35] scores and retain the top 5 hints<sup>4</sup>. Convergence measures how well the hints narrow down the space of possible answers. In other words, it reflects how effectively the hints guide a user toward eliminating incorrect answer candidates and focusing on the correct answer. This ensures that we retain only the most informative and accurate hints to construct high-quality passages. After this filtering and selection stage, we obtain a total of 17,203 questions paired with 81,235 high-convergence hints.

**3.1.3 Question Type and Difficulty.** To further enrich the dataset with metadata, we detect the type and estimate the difficulty of each question. For question type detection, we use the HintEval

<sup>4</sup>We select the top 5 hints to maintain computational feasibility, as generating and evaluating all subsets and permutations becomes prohibitively expensive beyond this number.

```

Question: <question>
Answer: <ground_truth>
Candidate: <candidate>
Is candidate correct? Choose between "Yes" or "No"

```

**Figure 3: Prompt of GPT-Eval. The placeholder <question> represents the question, <ground\_truth> indicates the correct answer, and <candidate> shows the answer generated by different LLMs.**

framework [36], which adopts the classification method proposed by Tayyar Madabushi and Lee [52]. To estimate question difficulty, we apply the Reference-based Question Complexity method [14], which assesses difficulty by analyzing how many of the question’s retrieved passages contain the correct answer and by measuring the semantic relevance between those passages and the question.

**3.1.4 Question Answering.** LLMs are capable of answering many questions directly from their parametric knowledge, without relying on external context [39]. This poses a challenge for our setup, as it would prevent us from properly evaluating the effect of our passages if an LLM can generate the correct answer solely from its parametric knowledge, without relying on any external context. To address this, we filter out questions that the LLMs used as Readers and those used for passage-labeling can correctly answer without any external context.

To reduce bias, we use models from different families and with varying parameter sizes. Specifically, we use Gemma 3 1B [53], Qwen 3 4B [58], and LLaMA 3.1 8B [13] for passage-labeling, and LLaMA 3.2 1B [13], Gemma 3 4B [53], and Qwen 3 8B [58] as Readers<sup>5</sup>. Each LLM is prompted with the question alone, without any supporting context. If at least one model produces the correct answer, we classify the question as *parametrically answerable*. Conversely, if none of the models provides the correct answer, the question is classified as *non-parametrically answerable*.

To compare generated answers with ground-truth answers, we use the GPT-Eval method [24], as lexical matching is insufficient in the LLM era. For example, an LLM may output *United States of America* while the gold answer is *USA*; a simple string matching would fail in this case, whereas GPT-Eval correctly recognizes their semantic equivalence. This method uses the prompt shown in Figure 3, which outputs *Yes* if the generated answer matches the ground-truth and *No* otherwise. For reliability, we repeat this evaluation three times per question to ensure that none of the selected LLMs can answer it parametrically.

Finally, we randomly sample 5,000 *parametrically answerable* questions as a temporary training set, since answer leakage<sup>6</sup> is less problematic during training. From the *non-parametrically answerable* pool, we sample 2,000 questions as a temporary test set, with the remaining questions forming the temporary development

<sup>5</sup>Other LLMs could also be used; however, we selected these models to create a balanced and fair evaluation environment.

<sup>6</sup>Answer leakage [36] measures how much a hint reveals the correct answer inside its content. It ensures that hints guide the user without directly revealing the solution.

**Table 1: Statistical Summary of the QUIT Dataset**

Metric	Value
Total Questions	7,401
Total Passages	2,405,325
Avg. Question Length (words)	12.56
Avg. Passage Length (words)	58.62
Avg. Answer Length (words)	1.90
Query-Passage Lexical Overlap	0.02
Answer Containment	0.00
Query-Passage Semantic Overlap	0.39

set. This guarantees that the development and test sets remain high-quality and unbiased for evaluation.

## 3.2 Dataset Preparation

In this stage, we generate passages for each question using their hints and label them to make the final QUIT benchmark.

**3.2.1 Subset and Permutation Generation.** From the previous stage, we obtain 8,095 questions and 40,475 hints across the training, development, and test sets. To construct diverse passages, we generate all possible non-empty subsets of the five selected hints for each question. For every subset, we then enumerate all possible permutations of its members. This results in a total of  $\sum_{k=1}^5 \binom{5}{k} k! = 325$  unique passages per question. For example, the subset containing all five hints yields 120 distinct passages, each expressing the same meaning but differing in sentence order.

**3.2.2 Passage Labeling.** We label all generated passages using LLMs. Salemi and Zamani [45] have demonstrated a strong correlation between passage relevance labels and the answerability of questions by LLMs. Building on this observation, for each question, we provide the question, a candidate passage, and the correct answer. We pass the question and the passage to three models: Gemma 3 1B, Qwen 3 4B, and LLaMA 3.1 8B to generate an answer based on the passage<sup>7</sup> using a few-shot prompting strategy<sup>8</sup> shown in Appendix A.4. To determine whether the answer produced by an LLM is correct, we use GPT-Eval, as described in Section 3.1.4. If at least one of the models generates the correct answer from the passage, we label it as **2 (fully relevant)**, since this indicates that the passage provides sufficient information for an LLM to infer the correct answer. In such cases, we also add the generated answer to the pool of valid answers for that question. Otherwise, we assign the passage the label **1 (partially relevant)**. The rationale is that even if a passage does not lead directly to the correct answer, it may still provide useful contextual information by indirectly describing the target entity. However, such passages often point to multiple candidate answers rather than uniquely identifying the correct one. For example, in Figure 1 (column *Answer-Supporting*), the *Relevant* passage enables the generation of *Lionel Messi*, whereas the *Partial* passage could lead to either *Lionel Messi* or *Diego Maradona*. Additionally, passages containing hints corresponding to other questions are labeled as **0 (irrelevant)** for the target question.

<sup>7</sup>Note that in Section 3.1.4 we evaluated questions without any context, whereas here we include the answer-supporting passages.

<sup>8</sup>We use a few-shot strategy to better guide the models toward producing valid answers.

**Table 2: Retriever performance on QUIT, MS MARCO, and Wikipedia corpora. We report Hit@k ( $k \in \{1, 10, 50, 100\}$ ) and MRR, with the best-performing retriever highlighted.**

Retriever	Corpus	Hit@1	Hit@10	Hit@50	Hit@100	MRR
BM25	QUIT	0.00%	0.25%	0.44%	0.57%	0.04%
	MS MARCO	23.09%	41.15%	56.77%	63.08%	29.31%
	Wikipedia	45.60%	72.22%	84.55%	88.25%	54.73%
DPR	QUIT	9.89%	16.62%	21.22%	23.74%	11.28%
	MS MARCO	14.41%	32.47%	49.19%	58.04%	20.55%
	Wikipedia	29.17%	57.41%	74.65%	80.32%	38.36%
ColBERT	QUIT	12.41%	16.44%	19.40%	20.21%	12.62%
	MS MARCO	33.28%	49.88%	62.79%	68.63%	38.94%
	Wikipedia	60.53%	83.04%	90.86%	92.94%	68.86%
Contriever	QUIT	6.49%	13.29%	18.95%	22.54%	8.15%
	MS MARCO	19.85%	43.63%	60.30%	66.72%	28.06%
	Wikipedia	45.37%	76.79%	87.44%	90.80%	56.56%
BGE	QUIT	12.85%	21.98%	27.96%	30.23%	14.68%
	MS MARCO	27.55%	47.69%	63.31%	71.93%	34.35%
	Wikipedia	53.47%	81.71%	91.72%	94.33%	63.31%

**3.2.3 Human Verification.** To ensure the quality and reliability of the development and test sets, we perform a human verification step. Specifically, we review the answers that GPT-Eval classified as correct, since LLMs—even advanced ones such as GPT—may occasionally produce incorrect or hallucinated outputs. Human validation thus helps us construct more trustworthy development and test sets. For this process, we designed an evaluation interface shown in Figure 7 in Appendix A.1 that enables annotators to verify the correctness of answers and resolve potential issues effectively.

**3.2.4 Label Updating.** During human verification of LLM-generated answers, new cases of answer leakage may inadvertently arise. For example, an answer marked as incorrect by GPT-Eval might later be judged correct by human annotators, and that new correct answer needs to be verified with respect to the risk of answer leakage. To address this issue, we reapply the *Question Filtering* procedure described in Section 3.1.2 to all questions and remove those for which at least one hint now exhibits answer leakage. Human verification may also affect passage labels. Specifically, if a passage no longer yields any correct answers after review, its label is updated from fully relevant (label 2) to partially relevant (label 1). After these adjustments, we obtain the cleaned and finalized versions of the training, development, and test sets, which together constitute the QUIT benchmark. Table 1 shows its statistics.

## 4 Experimental Setup

In our experiments, we evaluate diverse retrievers, rerankers, and readers on the Inferential QA task and QUIT benchmark.


For the *retriever*, we employ five retrieval methods: BM25 [43], DPR [25], ColBERT v2 [46], Contriever [19], and BGE [8], all accessed via the Pyserini toolkit [29]. We fine-tune DPR using the Tevatron library [15], and ColBERT using its official repository<sup>9</sup>.

For the *reranker*, we experiment with five models: LiT5 (T5-3B) [51], MonoT5 (T5-3B) [37], RankGPT (LLaMA 3.1 8B) [49], RankT5 (T5-3B) [62], and UPR (T0-3B) [44], implemented using the Rankify [1]. MonoT5 is fine-tuned using the PyGaggle [40].

<sup>9</sup><https://github.com/stanford-futuredata/ColBERT>

**Table 3: Comparison of rerankers across QUIT, MS MARCO, and Wikipedia corpora. We report Hit@k ( $k = 1, 10, 50$ ) and MRR. The best reranker for each corpus is highlighted, while the highest score is bolded.**

Retriever	Corpus	Reranker	Hit@1	Hit@10	Hit@50	MRR
BGE	QUIT	LiT5	26.03%	29.21%	33.80%	28.10%
		<b>MonoT5</b>	<b>27.60%</b>	<b>29.98%</b>	<b>32.35%</b>	<b>28.54%</b>
		RankGPT	24.02%	29.05%	33.33%	25.70%
		RankT5	26.62%	<b>30.44%</b>	32.52%	27.80%
		UPR	26.85%	29.86%	32.70%	27.89%
ColBERT	MS MARCO	LiT5	<b>40.28%</b>	53.53%	64.24%	44.80%
		MonoT5	38.60%	55.44%	65.10%	44.27%
		RankGPT	28.88%	49.88%	62.79%	36.76%
		<b>RankT5</b>	<b>39.53%</b>	<b>55.56%</b>	<b>65.39%</b>	<b>45.01%</b>
		UPR	36.40%	54.69%	65.22%	42.62%
ColBERT	Wikipedia	LiT5	68.52%	85.82%	91.38%	74.59%
		MonoT5	68.17%	88.02%	91.84%	75.33%
		RankGPT	52.60%	82.99%	90.86%	64.60%
		<b>RankT5</b>	<b>71.88%</b>	<b>88.48%</b>	<b>91.96%</b>	<b>78.06%</b>
		UPR	65.57%	86.52%	<b>92.13%</b>	73.37%



**Figure 4: Left: Hit@k comparison across different corpora. Right: Hit@100 with respect to the number of relevant passages (label 2) across corpora.**

For the *reader*, we employ three LLMs from different model families and parameter scales: LLaMA 3.2 1B [13] and Gemma 3 4B [53] as general-purpose LLMs, and Qwen 3 8B [58] as a reasoning-oriented model. This setup allows us to evaluate the effectiveness of LLMs with different specializations on inferential questions. As noted in Section 3.1.4, we choose these models because we verified that none of them can answer the questions solely from its parametric knowledge, which ensures a fair evaluation on passages. All models are executed on two NVIDIA A40 GPUs with 40GB.

We evaluate our experiments using standard metrics: **Hit@k**, which measures whether a gold passage appears in the top-k (reported for  $k = 1, 5, 10, 50, 100$ ); **Recall@k**, which evaluates the coverage of relevant passages (for  $k = 5, 10, 50$ ); **MRR**, the mean reciprocal rank of the first relevant passage; **NDCG@k**, which captures ranking quality with graded relevance (for  $k = 10$  and  $k = 100$ ); and **EM**, the exact match between predicted and gold answers in the reader component.

## 5 Results and Discussion

### 5.1 Compared Benchmarks

To evaluate how challenging the Inferential QA task is for current retrievers and rerankers, we compare QUIT corpus<sup>10</sup> against

<sup>10</sup>We refer to the union of all of answer-supporting passages from the training, development, and test sets as the QUIT corpus.

**Table 4: Retriever performance on vanilla and finetuned models. We report Hit@k, Recall@k, MRR, and nDCG@k.**

Retriever	Hit@K		Recall@K			MRR@K	nDCG@K	
	1	5	5	10	50	100	10	100
<b>Vanilla</b>								
BGE	23.73%	27.37%	0.37%	0.75%	3.80%	25.45%	18.95%	21.14%
DPR	18.58%	21.18%	0.28%	0.54%	2.54%	19.77%	13.68%	13.64%
ColBERT	20.31%	21.12%	0.31%	0.63%	3.15%	20.69%	16.14%	17.51%
<b>Finetuned</b>								
FT-DPR	20.91%	28.07%	0.32%	0.63%	3.12%	23.56%	14.98%	16.69%
FT-ColBERT	21.11%	23.84%	0.27%	0.53%	2.56%	20.83%	13.00%	13.98%

two widely used benchmarks: Wikipedia [25] and MS MARCO [3]. Specifically, we retrieve passages from all three corpora (Q<sub>UIT</sub>, Wikipedia, and MS MARCO) for the questions of the test set.

In traditional QA systems, relevance is typically defined by answer containment, i.e., a passage is considered relevant if it includes the correct answer. The *Answer-Containing* column in Figure 1 illustrates this definition. However, as explained in Section 3.2.2, our labeling method differs because our passages are designed to avoid answer leakage, making answer containment unsuitable. To address this, we relabel the retrieved passages from Wikipedia and MS MARCO using our method. Since partially relevant (label 1) in our framework is defined based on the relationship between a hint and a question, it cannot be directly applied to these corpora, which lack hints, we therefore use passages labeled as fully relevant (label 2), which are determined based on whether an LLM can correctly answer the question using the passage, to ensure a fair comparison with answer-supporting passages.

For fairness, we retrieve the top 100 passages from Q<sub>UIT</sub>, MS MARCO, and Wikipedia using five different retrievers, and the results are reported in Table 2. These results reveal a large gap in Hit@k and MRR between Q<sub>UIT</sub> and the other corpora, highlighting the inherent difficulty of retrieving relevant passages in the Inferential QA task. The table also shows that the best-performing retriever for Q<sub>UIT</sub> differs from those for MS MARCO and Wikipedia, further underscoring the difference between answer-supporting and answer-containing passages. Figure 4 further illustrates this gap. The left chart shows that Hit@k is consistently highest for Wikipedia, followed by MS MARCO, and lowest for Q<sub>UIT</sub>, demonstrating the increased challenge of retrieving from Q<sub>UIT</sub>. The right chart shows that although the retrieved passages of Q<sub>UIT</sub> include 289,660 relevant passages (label 2), Hit@100 is only 30.23%. In contrast, for MS MARCO and Wikipedia, which contain 31,195 and 52,945 relevant retrieved passages, Hit@100 reaches 68.63% and 92.94%, respectively. This indicates that, although one might expect a higher number of relevant passages to yield a higher Hit@100, this is not the case for Q<sub>UIT</sub>. Finally, it is worth noting that the average passage length in MS MARCO is 56.57 tokens, while in Q<sub>UIT</sub> it is 58.62 tokens, confirming that the observed gap is not due to differences in passage length.

In addition to retrievers, Table 3 shows that the best reranker for Q<sub>UIT</sub> differs from those for MS MARCO and Wikipedia, indicating that answer-supporting passages behave differently from answer-containing passages and that the prevailing hypothesis about what constitutes the most relevant passages does not hold for answer-supporting passages.

**Table 5: Performance comparison of rerankers applied to BGE and finetuned retrievers. We report Hit@K, Recall@K, MRR, and nDCG@K. The best reranker for each retriever is highlighted, and the highest score in each column is bolded.**

Retriever	Reranker	Hit@K		Recall@K			MRR@K	nDCG@K	
		1	5	5	10	50	100	10	100
BGE	LiT5	27.03%	29.17%	0.41%	0.82%	3.86%	28.10%	21.15%	21.49%
	MonoT5	<b>27.60%</b>	<b>29.46%</b>	<b>0.42%</b>	<b>0.84%</b>	<b>4.01%</b>	<b>28.54%</b>	<b>22.36%</b>	<b>21.81%</b>
	RankGPT	24.02%	27.49%	0.37%	0.75%	3.78%	25.70%	18.98%	21.14%
	RankT5	26.62%	29.05%	0.41%	0.82%	4.00%	27.80%	21.67%	21.69%
	UPR	26.85%	28.94%	0.42%	0.83%	4.01%	27.89%	21.06%	21.61%
FT-DPR	LiT5	27.84%	32.58%	0.40%	0.77%	3.08%	30.34%	19.31%	16.20%
	MonoT5	<b>28.01%</b>	<b>31.89%</b>	<b>0.40%</b>	<b>0.78%</b>	<b>3.39%</b>	<b>30.24%</b>	<b>20.34%</b>	<b>16.63%</b>
	RankGPT	20.25%	28.41%	0.31%	0.61%	2.90%	24.55%	14.21%	15.36%
	RankT5	26.97%	31.71%	0.40%	0.78%	3.39%	29.60%	19.98%	16.56%
	UPR	27.89%	<b>32.64%</b>	<b>0.41%</b>	0.78%	3.37%	<b>30.38%</b>	19.60%	16.49%
FT-ColBERT	LiT5	<b>23.32%</b>	25.41%	0.33%	0.63%	2.66%	24.53%	16.37%	14.52%
	MonoT5	22.69%	25.58%	<b>0.33%</b>	<b>0.65%</b>	<b>2.86%</b>	24.16%	<b>17.07%</b>	<b>14.80%</b>
	RankGPT	18.00%	23.84%	0.27%	0.53%	2.55%	20.75%	13.05%	13.98%
	RankT5	22.92%	25.52%	0.33%	0.64%	2.84%	<b>24.29%</b>	16.82%	14.74%
	UPR	22.80%	<b>25.81%</b>	0.33%	0.63%	2.84%	24.33%	16.31%	14.67%

**Table 6: Performance of retrievers with the finetuned reranker (FT-MonoT5). We report Hit@k, Recall@k, MRR, and nDCG@k for both vanilla and finetuned retrievers.**

Retriever	Hit@K		Recall@K			MRR@K	nDCG@K	
	1	5	5	10	50	100	10	100
<b>Vanilla</b>								
BM25	0.00%	0.06%	0.00%	0.00%	0.00%	0.05%	0.02%	0.01%
Contriever	12.69%	16.53%	0.18%	0.34%	1.53%	14.69%	8.95%	8.27%
BGE	23.44%	26.98%	0.37%	0.74%	3.76%	25.10%	18.67%	20.77%
DPR	18.29%	20.88%	0.28%	0.53%	2.50%	19.47%	13.50%	13.37%
ColBERT	20.07%	20.84%	0.31%	0.62%	3.11%	20.43%	15.89%	17.21%
<b>Finetuned</b>								
FT-DPR	19.91%	28.07%	0.30%	0.60%	2.93%	24.07%	13.96%	15.36%
FT-ColBERT	18.11%	23.84%	0.27%	0.53%	2.56%	20.83%	13.00%	13.98%

## 5.2 Retriever

To evaluate the performance of retrievers for the Inferential QA task with respect to both partially relevant (label 1) and fully relevant (label 2) passages, we test BGE (the best-performing retriever for Q<sub>UIT</sub>, as shown in Table 2) alongside DPR and ColBERT, each in both vanilla and fine-tuned variants. This setup allows us to analyze the impact of fine-tuning on retrieving answer-supporting passages.


For fine-tuning DPR and ColBERT, we experiment with various numbers of positive and negative training samples (1, 5, 10, 50, 100, and 200), inspired by the findings of Chang [7], who showed that increasing the number of positive samples during training can enhance retrieval performance. We find that the optimal configuration is 10 positives (and 10 negatives<sup>11</sup>) for DPR, and 50 positives for ColBERT. Full results across all configurations are reported in Table 10 and Table 11 in Appendix A.2.

As shown in Table 4, BGE consistently achieves the highest performance. Fine-tuning leads to a slight improvement for DPR but provides no noticeable gain for ColBERT. We hypothesize that this difference stems from architectural design: DPR employs two separate BERT encoders for questions and passages, which may help it better capture passages. In contrast, ColBERT relies on a single shared encoder with late interaction, which may limit its ability to adapt to the challenges of the Inferential QA task. Also, BGE’s

<sup>11</sup>We use the same number of positive and negative passages, and negatives are sampled from the positive passages of other questions.

**Table 7: Performance of rerankers under the oracle setting. We report nDCG@k ( $k = 5, 10, 50, 100$ ) for both vanilla and finetuned rerankers. The best results are highlighted.**

Reranker	nDCG@K			
	5	10	50	100
<b>Vanilla</b>				
LiT5	72.94%	75.49%	79.34%	82.99%
RankGPT	65.02%	69.74%	78.09%	82.24%
RankT5	78.96%	80.18%	84.69%	87.49%
UPR	78.56%	79.72%	84.30%	87.25%
MonoT5	82.01%	82.95%	86.46%	88.71%
<b>Finetuned</b>				
FT-MonoT5	83.56%	84.24%	87.08%	89.17%



**Figure 5: Trend of EM across different settings using DPR (vanilla and fine-tuned) as the retriever and MonoT5 (vanilla and fine-tuned) as the reranker, evaluated with LLaMA 3.2 1B, Gemma 3 4B, and Qwen 3 8B as readers in the RAG setup.**

strong performance may be attributed to its LLM-based architecture, which is better equipped to model deeper semantic connections and implicit clues—key characteristics of answer-supporting passages.

Overall, these results suggest that *fine-tuning offers only marginal benefits and is insufficient for tackling the complexity of the Inferential QA task. Addressing this challenge may require fundamentally new retrieval approaches and paradigms.*

### 5.3 Reranker

To evaluate the effect of rerankers on the Inferential QA task, we test the QUIT benchmark using five different rerankers. These rerankers are applied on top of BGE (the best retriever) as well as the fine-tuned versions of DPR and ColBERT, allowing us to analyze their impact on this task. Table 5 shows that MonoT5 achieves the best performance with BGE and is also the top reranker for FT-DPR and FT-ColBERT. Compared to the retriever results in Table 4, rerankers provide only marginal improvements, indicating that reranking alone is not sufficient to reliably promote passages to the top of the retrieved list. *This suggests that understanding answer-supporting passages is also challenging for rerankers.*

In addition to the vanilla rerankers, we fine-tune MonoT5 under different settings, varying the number of positive and negative samples as in Section 5.2. We find that the best configuration uses 10 positive passages, with detailed results provided in Table 12 in Appendix A.3. However, as shown in Table 6, the fine-tuned MonoT5 performs worse than its vanilla counterpart (Table 5), indicating that fine-tuning does not help rerankers adapt to the Inferential QA

**Table 8: LLM-based Reader results across different retriever-reranker strategies. We report Exact Match (EM) for three LLMs. UN refers to Union<sub>norm</sub> and UF to Union<sub>freq</sub> methods.**

Retriever	Reranker	K	Strategy	LLaMA-32-1b	Gemma-3-4b	Qwen-3-8b
<b>Optimal</b>						
Oracle	Oracle	Oracle	Oracle	40.68%	90.16%	62.50%
<b>Oracle Retriever + Reranker (Vanilla &amp; Finetuned)</b>						
Oracle	MonoT5	5	UF	20.25%	50.41%	34.32%
Oracle	FT-MonoT5	5	UN	21.12%	51.10%	34.43%
<b>Retriever (Vanilla &amp; Finetuned)</b>						
BGE	–	5	UF	4.11%	13.14%	9.61%
DPR	–	5	UN	3.24%	9.38%	7.29%
FT-DPR	–	5	UN	2.60%	8.22%	6.31%
ColBERT	–	5	UN	3.82%	11.86%	8.28%
FT-ColBERT	–	5	UF	3.36%	8.85%	6.42%
<b>Retriever (Vanilla &amp; Finetuned) + Reranker (Vanilla)</b>						
BGE	MonoT5	5	UN	4.98%	15.34%	12.38%
DPR	MonoT5	5	UN	4.22%	11.75%	9.20%
FT-DPR	MonoT5	3	UF	4.46%	12.50%	9.90%
ColBERT	MonoT5	3	UF	4.46%	11.98%	8.46%
FT-ColBERT	MonoT5	3	UF	4.22%	11.75%	8.45%
<b>Retriever (Vanilla &amp; Finetuned) + Reranker (Finetuned)</b>						
BGE	FT-MonoT5	5	UF	4.86%	13.89%	10.82%
DPR	FT-MonoT5	3	UF	4.10%	10.94%	8.75%
FT-DPR	FT-MonoT5	3	UN	4.17%	12.44%	8.80%
ColBERT	FT-MonoT5	5	UF	4.46%	11.81%	8.68%
FT-ColBERT	FT-MonoT5	5	UN	4.22%	11.69%	7.87%

task. Overall, *similar to retrievers, rerankers require new approaches and paradigms to effectively handle answer-supporting passages.*


We also evaluate rerankers under an oracle setting, where we assume the retriever is perfect and can retrieve all relevant passages. In this setup, we compare the vanilla and fine-tuned versions of MonoT5. Table 7 shows that fine-tuned MonoT5 achieves the best performance when retrievers are assumed to be oracle. This suggests that fine-tuning can help rerankers in principle, but the improvements are still small, even under ideal retrieval. Thus, the poor performance of fine-tuned MonoT5 in Table 6 may be attributed to the limitations of the retrievers rather than rerankers themselves. Nonetheless, the results confirm that fine-tuning alone is insufficient, and more advanced approaches are needed to handle answer-supporting passages effectively.

### 5.4 Reader

After retrieving and reranking passages, we use the top-1, top-3, and top-5 passages as input to the reader in the RAG setting. This setup allows us to evaluate retrievers and rerankers on the Inferential QA task. A key challenge in the top-3 and top-5 settings is that simply concatenating passages to form the context introduces redundancy, since multiple passages may contain overlapping hints<sup>12</sup>. To address this, we propose two methods for constructing the final context: (1) Union<sub>norm</sub> and (2) Union<sub>freq</sub>.

We represent each passage  $p_i = \{s_{i1}, \dots, s_{ij}\}$ , where  $s_{ij}$  denotes the  $j$ -th sentence of passage  $p_i$ . The Union<sub>norm</sub> method applies a standard union while preserving order, producing a new set of sentences that are concatenated into a passage. The Union<sub>freq</sub> method instead scores each sentence based on two factors: (i) the rank of the passage it belongs to ( $rank(P_i)$ ), and (ii) the position of the sentence within that passage ( $pos_i(s)$ ). The scoring function is defined as:

<sup>12</sup>This overlap occurs because, as discussed in Section 3.2, relevant passages for a question are generated from the subsets and permutations of the five hints designed for that question.



**Figure 6: Distribution of passages by the number of hints across three settings: Retriever without Reranker, Oracle Retriever with Reranker, and Optimal.**

$$\text{score}(s) = \alpha \cdot \sum_{i:s \in P_i} \frac{1}{\text{rank}(P_i)} + \beta \cdot \sum_{i:s \in P_i} \frac{1}{\text{pos}_i(s)} \quad (1)$$

where  $\alpha$  controls the weight given to higher-ranked passages and  $\beta$  controls the weight given to earlier sentences within each passage. In our experiments, we set  $\alpha = 0.6$  and  $\beta = 0.4$ <sup>13</sup>. After computing the scores, sentences are ranked in descending score order, and the top- $k$  are concatenated to form a passage.

Table 8 reports the results of different retriever and reranker setups in LLM-based Reader. The *Optimal* setting corresponds to an oracle scenario where both retriever and reranker perform perfectly. To simulate this setting, we exhaustively evaluate all fully and partially relevant passages for each question and mark a question as answerable if at least one of these passages enables the LLM to produce the correct answer, showing the upper bound achievable when using LLaMA 3.2 1B, Gemma 3 4B, and Qwen 3 8B as readers. These results demonstrate that if retrieval and reranking were perfectly solved, EM could reach as high as 90.16% with Gemma 3 4B, and similarly high values for the other LLMs. This indicates that QUIT is of high quality and can effectively support LLMs in generating correct answers.

However, under the *Oracle retriever* setting—where we assume perfect retrieval but rely on rerankers to order passages—the EM is only about half of the *Optimal* score, showing that current rerankers are not capable of ordering passages effectively. Fine-tuning also fails to close this gap, indicating that rerankers require new methods and paradigms tailored to the Inferential QA task. Moreover, retriever-only results are the weakest, far below the *Optimal* scores, highlighting that current retrievers are also inadequate and must move beyond existing paradigms. While rerankers provide slight improvements, they cannot resolve the fundamental challenges. As shown in Figure 5, fine-tuning retrievers can even degrade performance, while rerankers provide only marginal improvements.

Table 4 shows that the fine-tuned DPR achieves better performance than DPR-vanilla; however, in the LLM-based Reader stage, its EM is lower than that of the vanilla model. A possible explanation is that fine-tuned retrievers may retrieve more relevant passages overall, but their ranking is suboptimal. In this case, rerankers can help improve their ordering, leading to better results than the vanilla model when rerankers are applied. Fine-tuned rerankers

perform slightly worse than their vanilla counterparts in the LLM-based Reader, as shown in Table 6, indicating that they provide little benefit. However, this figure shows that combining a fine-tuned retriever with a fine-tuned reranker yields better performance than other setups, although the improvement is very small.

The results show that general-purpose LLMs, such as Gemma 3 4B, outperform reasoning-oriented LLMs like Qwen 3 8B, even with fewer parameters. This indicates that relying on reasoning-oriented LLMs is not an effective solution for the Inferential QA task. Overall, *the results confirm that current retrievers and rerankers are not well-suited for the Inferential QA task, and that fine-tuning alone is not a viable solution; entirely new approaches are needed.*

Figure 6 illustrates this challenge. Retrievers and rerankers perform best when identifying passages with three sentences, whereas the *Optimal* results show that the most helpful passages for LLM-based Reader contain five hints. This is because longer passages provide more information about the answer, enabling the LLM-based Reader to answer with higher confidence. The weaker performance of retrievers and rerankers on passages with four or five hints may explain, in part, why they struggle with passages.

## 6 Conclusion

In this paper, we introduced **Inferential QA**—a new reasoning QA task that challenges Question Answering (QA) systems to infer answers from *answer-supporting* passages rather than extract them verbatim. To enable this study, we constructed **QUIT** (QUESTions requiring Inference from Texts), a large-scale dataset of 7,401 questions and 2.4 million passages derived from high-convergence hints, labeled through a combination of LLM-based answerability and human verification. Our extensive evaluation across multiple retrievers, rerankers, and LLM-based readers reveals that existing QA methods—effective on traditional answer-containing datasets—struggle under inferential conditions. Retrievers often fail to identify the correct answer-supporting passages, rerankers provide only minor improvements, and fine-tuning leads to limited or inconsistent gains. Even reasoning-oriented LLMs do not outperform smaller general-purpose models, suggesting that inference-based reasoning remains a major bottleneck in current QA pipelines. These findings highlight a fundamental gap between extraction and inference in QA. By shifting focus toward questions that require reasoning over indirect clues, **Inferential QA** opens a new research direction for developing retrieval, reranking, and reasoning methods capable of drawing conclusions from subtle, distributed evidence. We believe this work paves the path towards future research on designing retrieval-augmented reasoning systems that more closely resemble human inferential comprehension.

## Acknowledgments

The computational results presented here have been achieved using the LEO HPC infrastructure of the University of Innsbruck. This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant #2143434, and in part by the Office of Naval Research contract #N000142412612. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

<sup>13</sup>Grid search on the development set indicated that  $\alpha = 0.6$  and  $\beta = 0.4$  achieve the highest EM.

## References

- [1] Abdelrahman Abdallah, Bhawna Piryani, Jamshid Mozafari, Mohammed Ali, and Adam Jatowt. 2025. Rankify: A comprehensive python toolkit for retrieval, re-ranking, and retrieval-augmented generation. *arXiv preprint arXiv:2502.02464* (2025).
- [2] Heba Abdel-Nabi, Arafat Awajan, and Mostafa Z. Ali. 2022. Deep learning-based question answering: a survey. *Knowl. Inf. Syst.* 65, 4 (Dec. 2022), 1399–1485. doi:10.1007/s10115-022-01783-5
- [3] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamee, Bhaskar Mitra, Tri Nguyen, et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016).
- [4] Amy E. Barth and Amy Elleman. 2017. Evaluating the Impact of a Multistrategy Inference Intervention for Middle-Grade Struggling Readers. *Language, Speech, and Hearing Services in Schools* 48, 1 (2017), 31–41. arXiv:https://pubs.asha.org/doi/pdf/10.1044/2016\_LSHSS-16-0041 doi:10.1044/2016\_LSHSS-16-0041
- [5] 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*, David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard (Eds.). Association for Computational Linguistics, Seattle, Washington, USA, 1533–1544. https://aclanthology.org/D13-1160/
- [6] Jannis Bulian, Christian Buck, Wojciech Gajewski, Benjamin Börschinger, and Tal Schuster. 2022. Tomayto, Tomahto. Beyond Token-level Answer Equivalence for Question Answering Evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 291–305. doi:10.18653/v1/2022.emnlp-main.20
- [7] Shuai Chang. 2025. Improving Dense Passage Retrieval with Multiple Positive Passages. *arXiv preprint arXiv:2508.09534* (2025).
- [8] Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216* (2024).
- [9] Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 2924–2936. doi:10.18653/v1/N19-1300
- [10] Debrup Das, Sam O’Nuallain, and Razieh Rahimi. 2025. RaDeR: Reasoning-aware Dense Retrieval Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 19970–19997. doi:10.18653/v1/2025.emnlp-main.1011
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. doi:10.18653/v1/N19-1423
- [12] Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. DROP: A Reading Comprehension Benchmark Requiring Discrete Reasoning Over Paragraphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 2368–2378. doi:10.18653/v1/N19-1246
- [13] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, et al. 2024. The Llama 3 Herd of Models. *arXiv e-prints*, Article arXiv:2407.21783 (July 2024), arXiv:2407.21783 pages. doi:10.48550/arXiv.2407.21783
- [14] Matteo Gabburo, Nicolaas Paul Jedema, Siddhant Garg, Leonardo F. R. Ribeiro, and Alessandro Moschitti. 2024. Measuring Retrieval Complexity in Question Answering Systems. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 14636–14650. doi:10.18653/v1/2024.findings-acl.872
- [15] Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2023. Tevatron: An Efficient and Flexible Toolkit for Neural Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (Taipei, Taiwan) (SIGIR '23)*. Association for Computing Machinery, New York, NY, USA, 3120–3124. doi:10.1145/3539618.3591805
- [16] Mor Geva, Daniel Khoshdel, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. *Transactions of the Association for Computational Linguistics* 9 (2021), 346–361. doi:10.1162/tacl\_a\_00370
- [17] Arthur C. Graesser and Natalie K. Person. 1994. Question Asking During Tutoring. *American Educational Research Journal* 31, 1 (1994), 104–137. arXiv:https://doi.org/10.3102/00028312031001104 doi:10.3102/00028312031001104
- [18] Helia Hashemi, Mohammad Aliannejadi, Hamed Zamani, and W. Bruce Croft. 2020. ANTIQUE: A Non-factoid Question Answering Benchmark. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II* (Lisbon, Portugal). Springer-Verlag, Berlin, Heidelberg, 166–173. doi:10.1007/978-3-030-45442-5\_21
- [19] 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* (2022). https://openreview.net/forum?id=jKN1pXi7b0
- [20] Anubhav Jangra, Jamshid Mozafari, Adam Jatowt, and Smaranda Muresan. 2025. Navigating the Landscape of Hint Generation Research: From the Past to the Future. *Transactions of the Association for Computational Linguistics* 13 (06 2025), 505–528. doi:10.1162/tacl\_a\_00751
- [21] Adam Jatowt, Calvin Gehrler, and Michael Färber. 2023. Automatic Hint Generation. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval (Taipei, Taiwan) (ICTIR '23)*. Association for Computing Machinery, New York, NY, USA, 117–123. doi:10.1145/3578337.3605119
- [22] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning. In *Second Conference on Language Modeling*. https://openreview.net/forum?id=Rwhi91ideu
- [23] 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)*, Regina Barzilay and Min-Yen Kan (Eds.). Association for Computational Linguistics, Vancouver, Canada, 1601–1611. doi:10.18653/v1/P17-1147
- [24] Ehsan Kamaloo, Nouha Dziri, Charles Clarke, and Davood Rafiei. 2023. Evaluating Open-Domain Question Answering in the Era of Large Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 5591–5606. doi:10.18653/v1/2023.acl-long.307
- [25] 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)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 6769–6781. doi:10.18653/v1/2020.emnlp-main.550
- [26] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics* 7 (2019), 452–466. doi:10.1162/tacl\_a\_00276
- [27] Hyo-Jung O. Lee, Hyeon-Jin Kim, and Myung-Gil Jang. 2005. Descriptive Question Answering in Encyclopedia. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, Masaaki Nagata and Ted Pedersen (Eds.). Association for Computational Linguistics, Ann Arbor, Michigan, 21–24. doi:10.3115/1225753.1225759
- [28] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 9459–9474. https://proceedings.neurips.cc/paper\_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf
- [29] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A Python Toolkit for Reproducible Information Retrieval Research with Sparse and Dense Representations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, Canada) (SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 2356–2362. doi:10.1145/3404835.3463238
- [30] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv e-prints*, Article arXiv:1907.11692 (July 2019), arXiv:1907.11692 pages. doi:10.48550/arXiv.1907.11692
- [31] Meixiu Long, Duolin Sun, Dan Yang, Junjie Wang, Yecheng Luo, Yue Shen, Jian Wang, Hualai Zhou, Chunxiao Guo, Peng Wei, et al. 2025. Diver: A multi-stage approach for reasoning-intensive information retrieval. *arXiv preprint arXiv:2508.07995* (2025).

- [32] Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2024. Large language models: A survey. *arXiv preprint arXiv:2402.06196* (2024).
- [33] Jamshid Mozafari, Abdelrahman Abdallah, Bhawna Piryani, and Adam Jatowt. 2024. Exploring Hint Generation Approaches for Open-Domain Question Answering. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 9327–9352. doi:10.18653/v1/2024.findings-emnlp.546
- [34] Jamshid Mozafari, Florian Gerhold, and Adam Jatowt. 2025. WikiHint: A Human-Annotated Dataset for Hint Ranking and Generation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Padua, Italy) (SIGIR '25). Association for Computing Machinery, New York, NY, USA, 3821–3831. doi:10.1145/3726302.3730284
- [35] Jamshid Mozafari, Anubhav Jangra, and Adam Jatowt. 2024. TriviaHG: A Dataset for Automatic Hint Generation from Factoid Questions. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 2060–2070. doi:10.1145/3626772.3657855
- [36] Jamshid Mozafari, Bhawna Piryani, Abdelrahman Abdallah, and Adam Jatowt. 2025. HintEval: A Comprehensive Framework for Hint Generation and Evaluation for Questions. *arXiv preprint arXiv:2502.00857* (2025).
- [37] Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document Ranking with a Pretrained Sequence-to-Sequence Model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 708–718. doi:10.18653/v1/2020.findings-emnlp.63
- [38] Hariom A Pandya and Brijesh S Bhatt. 2021. Question answering survey: Directions, challenges, datasets, evaluation matrices. *arXiv preprint arXiv:2112.03572* (2021).
- [39] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2463–2473. doi:10.18653/v1/D19-1250
- [40] Ronak Pradeep, Haonan Chen, Lingwei Gu, Manveer Singh Tamber, and Jimmy Lin. 2023. PyGaggle: A Gaggle of Resources for Open-Domain Question Answering. In *Advances in Information Retrieval*, Jaap Kamps, Lorraine Goerriot, Fabio Crestani, Maria Maistro, Hideo Joho, Brian Davis, Cathal Gurrin, Udo Kruschwitz, and Annalina Caputo (Eds.). Springer Nature Switzerland, Cham, 148–162.
- [41] Peng Qi, Haejun Lee, Tg Sido, and Christopher Manning. 2021. Answering Open-Domain Questions of Varying Reasoning Steps from Text. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 3599–3614. doi:10.18653/v1/2021.emnlp-main.292
- [42] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Jian Su, Kevin Duh, and Xavier Carreras (Eds.). Association for Computational Linguistics, Austin, Texas, 2383–2392. doi:10.18653/v1/D16-1264
- [43] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. 3, 4 (April 2009), 333–389. doi:10.1561/1500000019
- [44] Devendra Sachan, Mike Lewis, Mandar Joshi, Armen Aghajanyan, Wen-tau Yih, Joelle Pineau, and Luke Zettlemoyer. 2022. Improving Passage Retrieval with Zero-Shot Question Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3781–3797. doi:10.18653/v1/2022.emnlp-main.249
- [45] Alireza Salemi and Hamed Zamani. 2024. Evaluating Retrieval Quality in Retrieval-Augmented Generation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 2395–2400. doi:10.1145/3626772.3657957
- [46] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 3715–3734. doi:10.18653/v1/2022.naacl-main.272
- [47] Rulin Shao, Rui Qiao, Varsha Kishore, Niklas Muennighoff, Xi Victoria Lin, Daniela Rus, Bryan Kian Hsiang Low, Sewon Min, Wen-tau Yih, Pang Wei Koh, et al. 2025. ReasonIR: Training Retrievers for Reasoning Tasks. *arXiv preprint arXiv:2504.20595* (2025).
- [48] Hongjin SU, Howard Yen, Mengzhou Xia, Weijia Shi, Niklas Muennighoff, Han Yu Wang, Liu Haisu, Quan Shi, Zachary S Siegel, Michael Tang, Ruoxi Sun, Jinsung Yoon, Sercan O Arik, Danqi Chen, and Tao Yu. 2025. BRIGHT: A Realistic and Challenging Benchmark for Reasoning-Intensive Retrieval. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=ykuc5q381b>
- [49] 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. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 14918–14937. doi:10.18653/v1/2023.emnlp-main.923
- [50] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4149–4158. doi:10.18653/v1/N19-1421
- [51] Manveer Singh Tamber, Ronak Pradeep, and Jimmy Lin. 2023. Scaling down, lifting up: Efficient zero-shot listwise reranking with seq2seq encoder-decoder models. *arXiv preprint arXiv:2312.16098* (2023).
- [52] Harish Tayyar Madabushi and Mark Lee. 2016. High Accuracy Rule-based Question Classification using Question Syntax and Semantics. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, Yuji Matsumoto and Rashmi Prasad (Eds.). The COLING 2016 Organizing Committee, Osaka, Japan, 1220–1230.
- [53] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, and Others. 2025. Gemma 3 Technical Report. arXiv:2503.19786 [cs.CL] <https://arxiv.org/abs/2503.19786>
- [54] Mengqiu Wang et al. 2006. A survey of answer extraction techniques in factoid question answering. *Computational Linguistics* 1, 1 (2006), 1–14.
- [55] Yu Wang, Vijay Srinivasan, and Jin Hongxia. 2022. A New Concept of Knowledge based Question Answering (KBQA) System for Multi-hop Reasoning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 4007–4017. doi:10.18653/v1/2022.naacl-main.294
- [56] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). [https://openreview.net/forum?id=VjQlMeSB\\_J](https://openreview.net/forum?id=VjQlMeSB_J)
- [57] Benfeng Xu, Chunxu Zhao, Wenbin Jiang, Pengfei Zhu, Songtai Dai, Chao Pang, Zhuo Sun, Shuohuan Wang, and Yu Sun. 2023. Retrieval-Augmented Domain Adaptation of Language Models. In *Proceedings of the 8th Workshop on Representation Learning for NLP (RePLNLP 2023)*, Burcu Can, Maximilian Mozes, Samuel Cahyawijaya, Naomi Saphra, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Chen Zhao, Isabelle Augenstein, Anna Rogers, Kyunghyun Cho, Edward Grefenstette, and Lena Voita (Eds.). Association for Computational Linguistics, Toronto, Canada, 54–64. doi:10.18653/v1/2023.repl4nlp-1.5
- [58] An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* (2025).
- [59] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium. doi:10.18653/v1/D18-1259
- [60] Zongmeng Zhang, Jinhua Zhu, Wengang Zhou, Xiang Qi, Peng Zhang, and Houqiang Li. 2024. BoolQuestions: Does Dense Retrieval Understand Boolean Logic in Language?. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 2767–2779. doi:10.18653/v1/2024.findings-emnlp.156
- [61] Minjun Zhu, Yixuan Weng, Shizhu He, Kang Liu, and Jun Zhao. 2022. Reason-ChainQA: Text-based Complex Question Answering with Explainable Evidence Chains. In *2022 China Automation Congress (CAC)*, 5431–5436. doi:10.1109/CAC57257.2022.10055048
- [62] Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and Michael Bendersky. 2023. RankT5: Fine-Tuning T5 for Text Ranking with Ranking Losses. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Taipei, Taiwan) (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 2308–2313. doi:10.1145/3539618.3592047

## A Appendix

### A.1 Human Verification



Figure 7: The human annotation interface. The *Question* section displays the question and its correct answer, while the *Candidate Answers* section lists all answers generated by LLMs across 325 passages. Annotators manually verify and select the correct answers.

Table 9: Demographic information of evaluators

Evaluator	Gender	Age	Education Level
1	Female	19	High School
2	Female	30	Bachelor’s Degree
3	Male	40	Master’s Degree
4	Male	32	PhD

To improve the reliability of the development and test sets, we perform human verification of answers marked as correct by GPT-Eval, mitigating potential errors or hallucinations from LLMs. Annotators use a custom interface shown in Figure 7 to validate answer correctness and resolve ambiguous cases.

Specifically, we display all generated answers that GPT-Eval marked as correct and highlight those that exactly match the ground-truth answers extracted from their original sources (TriviaHG [35] and WikiHint [34]). Human annotators are then asked to review these answers and select or deselect the ones they believe have been correctly or incorrectly labeled. For each question, annotators see the question text, the generated answers, and the associated ground-truth answer(s). They are instructed to assess whether the generated answers are semantically equivalent to the ground-truth answer, even if they differ lexically (e.g., “USA” vs. “United States of America”). If a generated answer is deemed incorrect despite being accepted by GPT-Eval, it is manually unmarked. Conversely, answers that are correct but not highlighted can be selected to ensure no valid answers are missed. This process helps correct both false positives and false negatives, resulting in higher-quality passage labels for downstream evaluation. Table 9 summarizes the demographic background of the human annotators involved in this verification process.

### A.2 Finetuned Retrievers

To fine-tune ColBERT and DPR, we experiment with different numbers of positive and negative samples. Given the complexity and reasoning requirements of answer-supporting passages, we hypothesize that even a single positive example might be sufficient for effective fine-tuning. To explore this, we test a range of values—1, 5, 10, 50, 100, and 200—for both positive and negative samples across ColBERT and DPR, aiming to identify the most optimal configuration.

The results are presented in Table 10 and Table 11, which report performance for each combination. Based on these results, we find that using 50 positive samples yields the best performance for ColBERT, while 10 positive samples perform best for DPR. We therefore adopt these configurations as the final fine-tuned versions of ColBERT and DPR used in our experiments.

### A.3 Finetuned Reranker

To fine-tune rerankers, we experiment with MonoT5, as it is a widely used reranking model and serves as a strong baseline. Similar to retrievers, we test different numbers of positive and negative samples—1, 5, 10, 50, 100, and 200—to analyze the effect of training data size on reranking answer-supporting passages.

The results are presented in Table 12, which reports performance for each configuration. Based on these results, we find that using 10 positive samples provides the best performance for MonoT5. We therefore adopt this configuration as the final fine-tuned version of MonoT5 used in our experiments.

Table 10: Performance of finetuned ColBERT retrievers with varying numbers of positive passages. We report Hit@k, Recall@k, MRR, and nDCG@k. The best are highlighted.

Retriever	Hit@K		Recall@K			MRR@K	nDCG@K	
	1	5	5	10	50	100	10	100
<b>Vanilla</b>								
ColBERT	20.31%	21.12%	0.31%	0.63%	3.15%	20.69%	16.14%	17.51%
<b>Finetuned</b>								
ColBERT (1 pos)	13.95%	17.88%	0.21%	0.42%	1.95%	15.93%	10.40%	10.36%
ColBERT (5 pos)	15.74%	20.14%	0.24%	0.46%	2.20%	17.93%	11.61%	11.93%
ColBERT (10 pos)	16.90%	20.37%	0.25%	0.49%	2.40%	18.64%	12.21%	13.20%
ColBERT (50 pos)	<b>21.11%</b>	<b>23.84%</b>	<b>0.27%</b>	<b>0.53%</b>	<b>2.56%</b>	<b>20.83%</b>	<b>13.00%</b>	<b>13.98%</b>
ColBERT (100 pos)	18.46%	22.92%	0.29%	0.56%	2.75%	20.56%	13.80%	14.87%
ColBERT (200 pos)	17.77%	23.32%	0.28%	0.55%	2.69%	20.35%	13.21%	14.58%

Table 11: Performance of finetuned DPR retrievers with varying numbers of positive passages. We report Hit@k, Recall@k, MRR, and nDCG@k. The best are highlighted.

Retriever	Hit@K		Recall@K			MRR@K	nDCG@K	
	1	5	5	10	50	100	10	100
<b>Vanilla</b>								
DPR	18.58%	21.18%	0.28%	0.54%	2.54%	19.77%	13.68%	13.64%
<b>Finetuned</b>								
DPR (1 pos)	19.97%	24.83%	0.31%	0.61%	2.97%	22.42%	14.69%	16.05%
DPR (5 pos)	20.83%	25.98%	0.31%	0.62%	3.08%	23.50%	14.61%	16.44%
DPR (10 pos)	<b>20.91%</b>	<b>28.07%</b>	<b>0.32%</b>	<b>0.63%</b>	<b>3.12%</b>	<b>23.56%</b>	<b>14.98%</b>	<b>16.69%</b>
DPR (50 pos)	20.31%	26.97%	0.31%	0.62%	3.05%	23.70%	14.49%	16.17%
DPR (100 pos)	20.83%	27.43%	0.31%	0.61%	3.03%	24.19%	14.14%	15.86%
DPR (200 pos)	20.06%	25.75%	0.30%	0.60%	2.93%	24.07%	13.96%	15.36%

**System:** You are an assistant that answers questions based on the provided context. You just answer questions with exact answers. You do not use sentences as the response.

**User:** Use the context to answer the question under conditions: 1. Answer should not be sentences. It should be some words. 2. Do not generate "sorry" or "I cannot ..." sentences; instead, use "NO ANSWER". 3. Do not generate explanations, reasoning, or full sentences—only provide the exact answer. 4. If the answer cannot be guessed from the context, respond only with "NO ANSWER".

**Context:** He was the 44th President of the United States. He served as President from 2009 to 2017. He was the first African-American President of the United States. He was a member of the Democratic Party. He was born on August 4, 1961 in Honolulu, Hawaii.

**Question:** Who won the Nobel Peace Prize in 2009?

**Assistant:** Barack Obama

**Context:** The capital city of this country is Paris. This country is located in northwestern Europe. This country has a long history and has played a significant role in international affairs. The official language of this country is French. The currency used in this country is the Euro.

**Question:** Édouard Daladier became Prime Minister of which country in 1933?

**Assistant:** France

**Context:** It's the coldest season of the year. It's the season when snow falls in many regions. It's the season when many people celebrate Christmas and New Year's Eve. It's the season when days are shorter and nights are longer. It's the season when many animals hibernate.

**Question:** If you have a 'Mahonia Japonica', in which season will it be in flower?

**Assistant:** Winter

**Context:** It is a team sport that originated in the United States. It is played with an oval-shaped ball. The objective of the game is to score points by advancing the ball into the opposing team's end zone. Points can be scored by carrying the ball across the opponent's goal line, throwing it to a teammate in the end zone, or kicking it through the opponent's goalposts. The game is divided into four quarters, each lasting 15 minutes.

**Question:** Which sport is played under the 'Harvard Rules'?

**Assistant:** AMERICAN FOOTBALL

**Context:** He was born on April 20, 1889, in Braunau am Inn, Austria. He was the leader of the Nazi Party. He became the chancellor of Germany in 1933. He took the title of Führer und Reichskanzler in 1934. He initiated World War II in Europe by invading Poland on September 1, 1939.

**Question:** Who was made an honorary citizen of Haslach, Austria, in 1938, an honour withdrawn in 2004?

**Assistant:** Adolf Hitler

**Context:** Its capital is Beijing. Its population is more than 1 billion. This country has a history spanning more than 3,000 years of continuous civilization.

**Question:** Which country borders 14 others and uses a single time zone across its vast territory?

**Assistant:** *China*

**Figure 8: Few-shot prompt provided to the language model during inference. Each example demonstrates how to use hint-based context to answer questions concisely with short responses (or NO ANSWER if the context is insufficient). The final pair shows the new query to be answered, with the model-generated answer shown in red for clarity.**

**Table 12: Performance of finetuned MonoT5 rerankers with varying numbers of positive passages. We report nDCG@k.**

Reranker	nDCG@K					
	1	5	10	20	50	100
<b>Vanilla</b>						
MonoT5	80.83%	82.01%	82.95%	84.24%	86.46%	88.71%
<b>Finetuned</b>						
MonoT5 (1 pos)	71.72%	76.07%	78.35%	80.61%	83.88%	86.78%
MonoT5 (5 pos)	81.91%	82.94%	83.61%	84.67%	86.65%	88.87%
MonoT5 (10 pos)	82.95%	83.56%	84.24%	85.30%	87.08%	89.17%
MonoT5 (50 pos)	75.42%	77.87%	79.49%	81.29%	84.09%	86.92%
MonoT5 (100 pos)	75.58%	77.45%	78.73%	80.47%	83.47%	86.43%
MonoT5 (200 pos)	75.12%	77.14%	78.08%	80.07%	82.71%	85.44%

#### A.4 Few-shot Prompt

Figure 8 illustrates the few-shot prompt used to evaluate language models. It comprises several *question–context–answer* examples, where each context is constructed by concatenating hints that indirectly describe an entity, concept, or event. The prompt demonstrates how to infer answers from contextual clues without generating explanations, guided by a system instruction that enforces short, phrase-level outputs, the use of "NO ANSWER" when information is insufficient, and the avoidance of reasoning or justification. The final example presents a new inferential question with unseen hints, prompting the model to produce its answer following the learned pattern.