

# BRIGHT: A REALISTIC AND CHALLENGING BENCHMARK FOR REASONING-INTENSIVE RETRIEVAL

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Existing retrieval benchmarks primarily consist of information-seeking queries (e.g., aggregated questions from search engines) where keyword or semantic-based retrieval is usually sufficient. However, many complex real-world queries require in-depth reasoning to identify relevant documents that go beyond surface form matching. For example, finding documentation for a coding question requires understanding the logic and syntax of the functions involved. To better benchmark retrieval on such challenging queries, we introduce **BRIGHT**, the first text retrieval benchmark that requires *intensive reasoning* to retrieve relevant documents. Our dataset consists of 1,384 real-world queries spanning diverse domains, such as economics, psychology, mathematics, and coding. These queries are drawn from naturally occurring and carefully curated human data. Extensive evaluation reveals that even state-of-the-art retrieval models perform poorly on **BRIGHT**. The leading model on the MTEB leaderboard (Muennighoff et al., 2023) SFR-Embedding-Mistral (Meng et al., 2024), which achieves a score of 59.0 nDCG@10,<sup>1</sup> produces a score of nDCG@10 of 18.3 on **BRIGHT**. We show that incorporating explicit reasoning about the query improves retrieval performance by up to 12.2 points. Moreover, incorporating retrieved documents from the top-performing retriever boosts question-answering performance by over 6.6 points. We believe that **BRIGHT** paves the way for future research on retrieval systems in more realistic and challenging settings.<sup>2</sup>

## 1 INTRODUCTION

Information retrieval is a widely employed technology that assists users in locating relevant information from extensive corpora, containing documents, web pages, and logging records (Bajaj et al., 2016; Thakur et al., 2021). Relevant information can relate to user queries in different ways—sometimes through straightforward matching patterns like shared keywords or semantic similarities, and other times through deeper, more nuanced connections such as analogous underlying principles. In many real-world scenarios, user queries can be highly complex, and finding the relevant documents requires intensive reasoning. For instance, an economist might want to find a story explained by the same economic theory as another story, or a programmer might want to use an error message to locate the corresponding syntax documentation. For these applications, relevant documents cannot be directly retrieved through lexical or semantic matching alone, but instead require additional reasoning steps to identify.

In this work, we study the problem of reasoning-intensive retrieval with **BRIGHT**, a new benchmark that requires intensive reasoning to retrieve relevant documents. Existing retrieval benchmarks, such as BEIR (Thakur et al., 2021) and MTEB (Muennighoff et al., 2023), primarily focus on fact-based queries typically derived from search engines, where the relevance between queries and documents is often straightforward and can be detected through simple keyword or semantic matching (Lee et al., 2019; Karpukhin et al., 2020). These datasets focus on retrieving specific pieces of information (e.g., “the widest highway in North America”), which leads to the relevant documents often having high lexical or semantic overlap with the queries. In contrast, the relevance between queries and documents in **BRIGHT** is not easily detectable through simple keyword or semantic matching, and requires deliberate reasoning due to the complex nature of our domains and queries (Figure 1).

<sup>1</sup>Retrieved from the MTEB leaderboard on 2024-05-28.

<sup>2</sup>Our code and data are available at <https://brightbenchmark.github.io>.

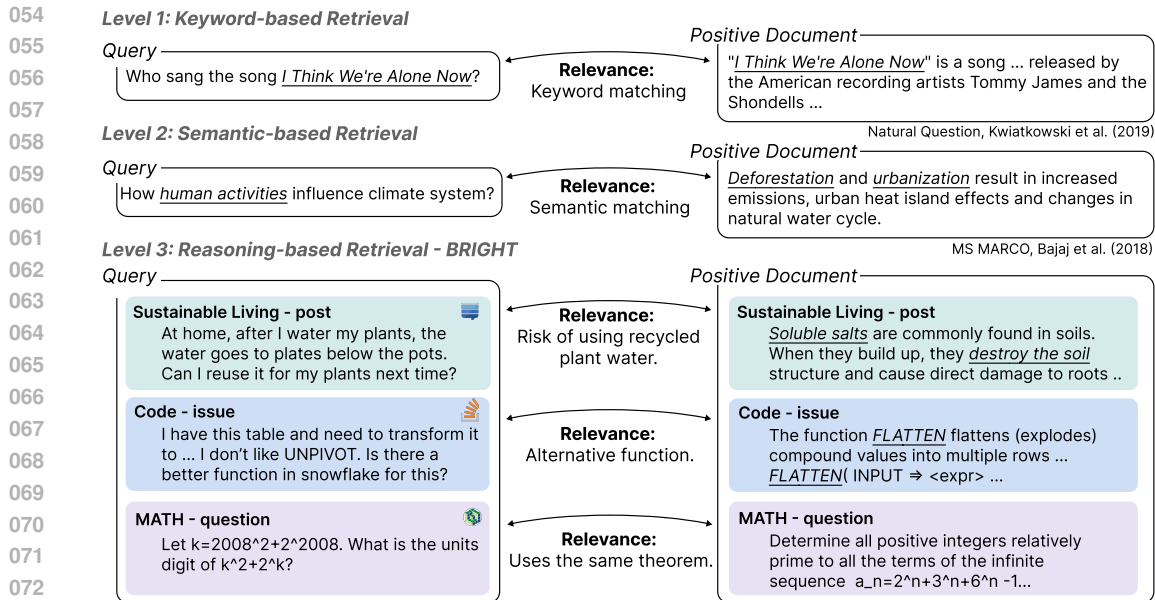


Figure 1: Existing retrieval benchmarks focus on keyword-based retrieval (level 1), or semantic-based retrieval (level 2), e.g., NQ, MS MARCO datasets (Kwiatkowski et al., 2019; Bajaj et al., 2018). **BRIGHT introduces level 3 retrieval, where the relevance between queries and documents requires intensive reasoning to determine.** Our data consists of natural user queries from diverse domains (e.g., economics, math, earth sciences, etc.). **BRIGHT** corpora also span across web data, such as blogs, syntax documentation, and STEM problem-solutions.

**BRIGHT** consists of 12 datasets from diverse and advanced domains, sourced from naturally occurring human data and meticulously curated sources. The benchmark comprises two main components: 1) seven datasets are constructed from StackExchange, where relevance is defined by whether a document is cited in the answer, and further validated through multiple annotators to ensure that the positive documents effectively support addressing queries. Given the inherent subjectivity of this assessment, we only include documents unanimously agreed upon by the annotators. 2) The remaining five datasets focus on coding and math problems, where the queries are inherently linked to the positive documents due to their shared underlying algorithms or theorems.

We evaluate 13 representative retrieval models of varying sizes and architectures. Comprehensive experiments highlight the challenges posed by **BRIGHT**, as the best-performing model, SFR-Embedding-Mistral (Meng et al., 2024), which scores 59.0 on the MTEB retrieval subset, BEIR (Thakur et al., 2021), achieves only an nDCG@10 score of 18.3 on **BRIGHT**. To identify promising directions for improving **BRIGHT**, we explore various strategies—one effective approach leverages LLMs to generate chain of thought reasoning steps (Wei et al., 2022) about the query before retrieval, resulting in average improvements of up to 12.2 points. Furthermore, augmenting the downstream model with retrieved documents from the top-performing retriever, Qwen, significantly enhances the question-answering performance by up to 6.6 points compared to the closed-book setting. However, using the oracle documents results in a 17.2-point improvement, highlighting the potential benefits of improving retriever models.

Moreover, our results demonstrate the robustness of **BRIGHT** against potential data leakage during large-scale pre-training, as no substantial performance gains are observed even when models are further trained on documents from the benchmark dataset. We hope that our findings inspire research directions to advance the state of the art in reasoning-intensive retrieval.

## 2 RELATED WORK

**Benchmarking retrieval.** Existing information retrieval (IR) datasets are typically constructed for information-seeking tasks, such as question-answering (Voorhees & Tice, 2000; Craswell et al., 2020;

Kwiatkowski et al., 2019; Chen et al., 2017; Maia et al., 2018), claim verification (Thorne et al., 2018; Diggelmann et al., 2020; Wadden et al., 2020), or entity retrieval (Hasibi et al., 2017; Petroni et al., 2021). Recent works expand retrieval benchmarks with more scenarios, such as instruction-following (Su et al., 2023; Weller et al., 2024; Oh et al., 2024; Li et al., 2024), multi-hop (Yang et al., 2018), and long-context retrieval (Saad-Falcon et al., 2024; Zhu et al., 2024). Comprehensive benchmarks like BEIR (Thakur et al., 2021) evaluate retrieval systems on diverse domains and tasks, with relevant documents sharing high semantic overlap with the query. Closest to our work, BIRCO (Wang et al., 2024) is designed to evaluate retrieval systems based on multifaceted objectives by leveraging existing datasets. However, it is limited to the LLM reranking setting and uses only a small candidate pool ( $\sim 100$  documents) for each query. RAR-b (Xiao et al., 2024) adapts existing commonsense, math, and code datasets into a retrieval setting to test whether models can directly retrieve answers to reasoning problems. However, we focus on a more realistic scenario where the answers are unlikely to be found as a substring in documents.<sup>3</sup> While both of the benchmarks focus on reasoning-intensive retrieval, **BRIGHT** is the first benchmark to collect realistic user queries and align them with relevant natural documents from large corpora, requiring deep reasoning to identify the correct matches.

**Dense retrieval models and retrieval augmented generation.** State-of-the-art retrieval systems often use dense models to encode text with rich representation. These models are trained on unsupervised data (Lee et al., 2019; Izcard et al., 2022), supervised data (Su et al., 2023; Asai et al., 2023a; Muennighoff, 2022), as well as LLM-generated data (Lee et al., 2024; Wang et al., 2023a; Muennighoff et al., 2024). In this work, we benchmark a diverse set of models across different axes: sparse and dense; small and large; open-source and proprietary. Additionally, as dense generative models continue to improve, retrieval-augmented generation (RAG; Asai et al., 2020; Borgeaud et al., 2022; Asai et al., 2023b; Muennighoff et al., 2024; Asai et al., 2024; Gao et al., 2023; Shi et al., 2024), which retrieves relevant documents to help generate coherent answers, has become an important application. Although we mainly focus on retrieval in this work, we conduct initial analyses and demonstrate that using stronger retrievers improves model generation for reasoning-intensive tasks.

**Benchmarking reasoning.** Many benchmarks aim to evaluate the reasoning abilities of LLMs, especially focused on mathematics and coding. As for mathematics, for example, datasets include GSM8K (Cobbe et al., 2021) and its extensions GSM1K (Zhang et al., 2024), TheoremQA (Chen et al., 2023a), MATH (Hendrycks et al., 2021), and LeanDojo (Yang et al., 2023). As for coding, HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), and LiveCodeBench (Jain et al., 2024) are often used. These benchmarks contain question-answer pairs and are usually sourced from textbooks, online resources, competitions, or domain experts. We source queries from selected high-quality datasets and construct **BRIGHT** through additional annotations, creating a realistic reasoning-intensive retrieval benchmark.

### 3 CONSTRUCTING BRIGHT

We introduce **BRIGHT**, a retrieval benchmark that tests whether retrieval systems can match queries and documents whose relevance requires intensive reasoning to solve, beyond just lexical and semantic similarities. In this section, we first formulate the task of reasoning-intensive retrieval (Section 3.1). Then, we detail the data collection process for the data from StackExchange (Section 3.2), coding datasets (Section 3.3), and theorem-based questions (Section 3.4). In Table 1, we present the benchmark statistics.

#### 3.1 TASK FORMULATION

Given a query  $Q$  and the retrieval corpus  $\mathcal{D} = \{D_1, \dots, D_n\}$ , retrievers are tasked to find relevant documents  $\mathcal{D}_Q^+ = \{D_{Q,1}^+, \dots, D_{Q,m}^+\} \subset \mathcal{D}$ , where  $m \ll n$  (positive). Negative documents are defined as  $\mathcal{D}_Q^- = \mathcal{D} \setminus \mathcal{D}_Q^+$ . In reasoning-intensive retrieval, the relevant document set  $\mathcal{D}_Q^+$  is connected to the query  $Q$  through specific reasoning traces or explanations (e.g., underlying principles, algorithms, or theorems) related to the query. For instance, common reasoning traces might involve identifying the query’s intent, analyzing and modeling the problem, and drawing sub-conclusions based on the provided descriptions. Such reasoning traces are typically absent from the query itself, making direct retrieval using only the query very challenging.

<sup>3</sup>We compare to RAR-b in detail in Appendix H.

Table 1: **Data statistics of BRIGHT.** For each dataset, we show the number of queries ( $\mathcal{Q}$ ) and documents ( $\mathcal{D}$ ), the average number of positive documents ( $\mathcal{D}^+$ ) per example, the average length of queries and documents (measured by the GPT-2 tokenizer Radford et al. (2019)), and sources of queries and documents. Q&Sol refers to demonstration examples of question-solution pairs. TheoremQA-Q and TheoremQA-T refer to question retrieval and theorem retrieval based on TheoremQA respectively. Examples for each dataset can be found in Appendix C.

Dataset	Total Number			Avg. Length		Source		Examples
	Q	$\mathcal{D}$	$\mathcal{D}^+$	Q	$\mathcal{D}$	Q	$\mathcal{D}$	
<i>StackExchange</i>								
Biology	103	57,359	3.6	115.2	83.6	StackExchange post	Web pages: article, tutorial, news, blog, report ...	Tab. 20
Earth Science	116	121,249	5.3	109.5	132.6			Tab. 21
Economics	103	50,220	8.0	181.5	120.2			Tab. 22
Psychology	101	52,835	7.3	149.6	118.2			Tab. 23
Robotics	101	61,961	5.5	818.9	121.0			Tab. 24
Stack Overflow	117	107,081	7.0	478.3	704.7			Tab. 25
Sustainable Living	108	60,792	5.6	148.5	107.9			Tab. 26
<i>Coding</i>								
LeetCode	142	413,932	1.8	497.5	482.6	Coding question	Coding Q&Sol	Tab. 27
Pony	112	7,894	22.5	102.6	98.3	Coding question	Syntax Doc	Tab. 28
<i>Theorems</i>								
AoPS	111	188,002	4.7	117.1	250.5	Math Olympiad Q	STEM Q&Sol	Tab. 29
TheoremQA-Q	194	188,002	3.2	93.4	250.5	Theorem-based Q	STEM Q&Sol	Tab. 30
TheoremQA-T	76	23,839	2.0	91.7	354.8	Theorem-based Q	Theorems	Tab. 31

### 3.2 STACKEXCHANGE: RETRIEVING WEB PAGES THAT HELP ANSWER QUESTIONS

**Relevance:** A document is considered relevant to a query only if it is cited in an accepted or highly voted answer and unanimously confirmed by annotators and domain experts that it helps reason through the query with critical concepts or theories.

StackExchange<sup>4</sup> is a popular community-driven platform where users ask questions and receive answers from other users. Among its 170+ sites, we select 7 diverse and knowledge-intensive domains: Sustainable Living, Economics, Psychology, Robotics, Earth Science, Biology, and coding in Stack Overflow. Unlike the short questions in traditional retrieval benchmarks, questions on StackExchange often contain long and technical descriptions of the problems and end with a logically complex question, such as fixing a bug. Responses often link to external web pages that contain relevant information to address the question. We construct query-document pairs based on user posts and documents referenced in the answers (see Figure 2).

**Selecting posts.** Human annotators<sup>5</sup> browse posts from newest to oldest and select a post with at least one answer that (1) is accepted by the user or receives  $> 5$  votes, and (2) contains one or more URL links. This process ensures that each dataset has a sufficient number of high-quality examples.

**Constructing query and positive documents.** For each selected post, we construct the query and positive documents as follows:

Step 1: The annotator combines the title and content of the post to form the query  $Q$ .

Step 2: The annotator visits web pages linked in the answers and includes them as positive documents if they are relevant to queries or discards them otherwise.

Step 3: If no web page is considered positive, the post itself is discarded. For each collected web page, the annotator splits the content into passages and selects positives  $\mathcal{D}_{Q,i}^+$  following the relevance definition.

**Constructing negative documents.** To prevent models from relying purely on lexical or semantic similarities, we ensure that the negative documents for each query also address similar topics.

<sup>4</sup><https://stackexchange.com/>

<sup>5</sup>authors and college students from corresponding fields.

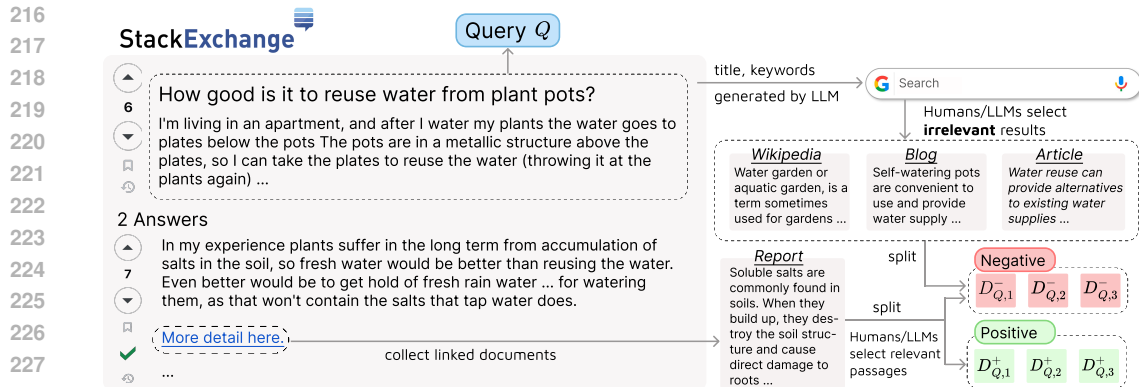


Figure 2: An overview of the data annotation process for StackExchange data. The post content is used as the query. Positive documents are selected passages from web pages linked in the answer, while the remaining passages, results from Google, and passages filtered by annotators are considered negatives. The web pages can include content from Wikipedia, blogs, articles, reports, and more.

Specifically, annotators gather search results from Google on the same topic and select documents that, while topically related, do not meet the specific requirements of the query (for more examples, see Appendix C). The collection procedure is as follows:

- Step 1: Annotators search Google using the posts’ title or LLM-summarized post keywords, and identify web pages that are semantically similar but not relevant to answering the question.
- Step 2: Annotators collect up to 5 negative web pages for each query and split them into hard negative passages, which consist of  $\mathcal{D}_Q^-$ .

For each dataset, we compile all the collected passages into a unified retrieval corpus  $\mathcal{D}$ . For any given query, all passages in the corpus—excluding its positive passages—are treated as negatives, including positive and negative passages associated with other queries. In contrast to traditional retrieval tasks such as open-domain QA (Fan et al., 2019; Kwiatkowski et al., 2019), where the retrieval pool typically includes documents that directly answer the query, we simulate a realistic scenario where positive documents only provide useful information to help users derive an answer. Please find more details on the construction of the dataset in Appendix B.1.

The annotation process involves a single computer science student performing the initial annotations, which are then verified by two PhD students specializing in the corresponding fields. Only annotations that receive unanimous approval from all involved parties, the original annotator and two expert reviewers, are retained in the final dataset. This approval process applies to both the positive and negative document annotations, ensuring a high standard of accuracy and reliability in the annotated examples. See more guidelines to annotators in Appendix E.

### 3.3 CODING: RETRIEVING DOCUMENTATION OR SIMILAR SOLVED PROBLEMS

**Relevance:** The relevance between queries and positive documents is defined by whether the coding problem (i.e., query) either requires the corresponding syntax documentation or involves the same algorithm and/or data structure.

To solve a coding problem, programmers often need to refer to the documentation or find similar problems that share the same algorithmic design. However, given only a problem description, it is difficult to find relevant documentation or similar problems via simple keyword or semantic matching. We construct two retrieval datasets on coding, where the relevance between queries and documents is grounded in the syntax usage and algorithm design.

**Pony.** When working with a rare programming language, consulting the manual can be invaluable for understanding its syntax and function usage. However, in such cases, the problem description would likely have low semantic similarity and lexical overlap with the relevant documentation. This discrepancy necessitates intensive reasoning in identifying the particular syntax or function that



is relevant to the problem at hand. We adopt a code generation dataset featuring Pony (Su et al., 2024), a rare programming language, and construct a retrieval dataset. We use the instructions of coding problems as queries  $Q$ , the annotated documentation about the required syntax as the positive documents  $\mathcal{D}_Q^+$ , and the complete language manual as the retrieval pool of documents  $\mathcal{D}$ , where each  $D_i$  contains descriptions about syntax usage of Pony, such as conditionals, loops, and classes.

**LeetCode.** We also explore coding problems that deal with algorithms and data structures, where the goal is to retrieve problems and solutions that share the same algorithmic design. We source coding problems and solutions from LeetCode.<sup>6</sup> The problem descriptions are used as queries  $Q$ , and the positive documents  $\mathcal{D}_Q^+$  are solved problems (with solutions) that were annotated as similar problems by LeetCode. Each document  $D_i = (Q_i, A_i)$  from LeetCode contains a problem statement  $Q_i$  and a Python solution  $A_i$ . To increase difficulty, we only keep questions that are grounded in real-world scenarios, where arriving at the key algorithm or data structure requires intensive reasoning. We construct a large corpus  $\mathcal{D}$  by combining questions and solutions from LeetCode and Python code from CodeSearchNet (Husain et al., 2019). See Appendix B.6 for more details on the dataset.

### 3.4 THEOREM-BASED QUESTIONS: RETRIEVE SOLVED PROBLEMS USING THE SAME THEOREMS OR RELEVANT THEOREM STATEMENTS

**Relevance:** A query (i.e., a solved problem) is relevant to a document if the document references the same theorem used in the query.

When tackling a new math or physics problem, users often reference previously solved problems or directly consult relevant theorem statements to guide their reasoning. Retrieving such similar problems or theorems can be challenging, as problems that share the same underlying logic may differ significantly in surface form, as shown in Table 30. In this setting, the query  $Q$  is a theorem-based question, and the corpus  $\mathcal{D}$  consists of solved STEM problems  $D_i = (Q_i, A_i)$ , where  $Q_i$  is the problem statement and  $A_i$  is its solution, or  $D_i$  are theorem documents from formal theorem collections such as ProofWiki. We consider  $D_i$  as a positive document if it shares the same theorem as the query’s solution. The dataset is built from high-quality STEM sources (Cobbe et al., 2021; Yuan et al., 2023; Hendrycks et al., 2021; Ling et al., 2017; Chen et al., 2023a; Li et al., 2023a), and for details on corpus construction, refer to Appendix B.2.

**TheoremQA.** Derived from textbooks, online resources, and experts, TheoremQA (Chen et al., 2023b) contains questions that are based on specific mathematical or scientific theorems (e.g., the binomial theorems), and represent problems that students and other users might encounter in their studies. To ensure that the model does not simply rely on the surface-level wording of the questions, we use GPT-4<sup>7</sup> to rephrase the question into more concrete, applied scenarios while maintaining the same required theorem. The prompts used for rewriting the questions and an example are shown in Table 13. Human annotators carefully review the rewritten questions and make necessary revisions to ensure that they are valid and consistent with the original questions. A document  $D_i = (Q_i, A_i)$  is positive if  $A_i$  uses the same theorem as the query’s solution. Additional details are in Appendix B.3.

**AoPS.** Math competition problems have been widely used to evaluate the problem-solving skills of students and LLMs (Hendrycks et al., 2021). Sourced from American and International Math Olympiads, these problems often require the application of advanced mathematical theorems and techniques, such as Fermat’s Little Theorem or Ptolemy’s theorem. To practice for the competitions, students often learn by solving other problems that require the same problem-solving skills. To this end, we collect a new dataset of math competition problems, called AoPS, annotated with their respective problem-solving skills from AoPS Wiki<sup>8</sup>. The collected problem-solving skills are shown in Table 18. Similar to TheoremQA, we consider a solved math problem  $D_i = (Q_i, A_i)$  positive if its solution uses the same problem-solving skill as the query’s solution. From preliminary qualitative analysis, we find that competition problems are deliberately written in diverse ways such that it is challenging to identify the required techniques; thus, we do not rephrase the problem statements. Details are in Appendix B.5.

<sup>6</sup><https://leetcode.com/>

<sup>7</sup>GPT-4 refers to the version gpt-4-0125-preview throughout this work.

<sup>8</sup><https://artofproblemsolving.com/wiki/>

Table 2: **The performance of retrieval models on BRIGHT.** We report nDCG@10 for all datasets: Biology (Bio.), Earth Science (Earth.), Economics (Econ.), Psychology (Psy.), Robotics (Rob.), Stack Overflow (Stack.), Sustainable Living (Sus.), LeetCode (Leet.), Pony, AoPS, TheoremQA with question retrieval (TheoQ.) and with theorem retrieval (TheoT.). Avg. denotes the average score across 12 datasets. The best score on each dataset is shown in bold and the second best is underlined. We show that reasoning-intensive retrieval is challenging for current retrievers, where the best model only achieves an nDCG@10 score of 24.3 on average. Model details are in Appendix A.1

	StackExchange						Coding		Theorem-based			Avg.	
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.	TheoT.	
<i>Sparse model</i>													
BM25	19.2	27.1	14.9	12.5	13.5	16.5	15.2	24.4	7.9	6.0	13.0	6.9	14.8
<i>Open-sourced models (&lt;1B)</i>													
BGE	12.0	24.2	16.6	17.4	12.2	9.5	13.3	26.7	5.6	6.0	13.0	6.9	13.6
Inst-L	15.6	21.5	16.0	21.9	11.5	11.2	13.2	20.0	1.3	8.1	20.9	9.1	14.2
SBERT	15.5	20.1	16.6	22.6	8.4	9.5	15.3	26.4	6.9	5.3	20.0	10.8	14.8
<i>Open-sourced models (&gt;1B)</i>													
E5	18.8	26.0	15.5	15.8	16.4	9.8	18.5	28.7	4.8	7.1	<u>26.1</u>	<u>26.8</u>	17.9
SFR	19.5	26.6	17.8	19.0	16.7	12.7	<u>19.8</u>	27.4	2.0	7.4	24.3	26.0	18.3
Inst-XL	21.9	<u>34.4</u>	<b>22.8</b>	27.4	<b>17.4</b>	<u>19.1</u>	18.8	27.5	5.0	8.5	15.6	5.9	18.7
GritLM	<u>25.0</u>	32.8	19.0	19.9	<u>17.3</u>	11.6	18.0	<u>29.8</u>	<b>22.0</b>	8.8	25.1	21.1	<u>20.9</u>
Qwen	<b>30.9</b>	<b>36.2</b>	17.7	24.6	13.5	<b>19.9</b>	14.9	25.5	<u>14.4</u>	<b>27.8</b>	<b>32.9</b>	<b>32.9</b>	<b>24.3</b>
<i>Proprietary models</i>													
Cohere	19.0	27.5	<u>20.2</u>	21.8	16.2	16.5	17.7	26.8	1.8	6.5	15.1	7.1	16.3
OpenAI	23.7	26.3	20.0	<u>27.5</u>	12.9	12.5	<b>20.3</b>	23.6	2.5	8.5	23.8	12.3	17.8
Voyage	23.6	25.1	19.8	24.8	11.2	15.0	15.6	<b>30.6</b>	1.5	7.4	<u>26.1</u>	11.1	17.7
Google	23.0	<u>34.4</u>	19.5	<b>27.9</b>	16.0	17.9	17.3	29.6	3.6	<u>9.3</u>	21.5	14.3	19.5

**Theorem retrieval.** Besides similar problems, retrieving relevant theorem definitions is also helpful. We use the queries from the aforementioned TheoremQA dataset with a different corpus  $\mathcal{D}$ , where each document  $D_i$  is a theorem statement from ProofWiki<sup>9</sup>. We align theorems in theoremQA with ProofWiki documents using title matching, followed by GPT-4 verification to ensure a candidate theorem is used in the solution. We retain only queries with at least one relevant theorem statement. Manual annotation of relevance between problems and documents also showed substantial agreement (Cohen’s  $\kappa = 0.62$ ) between human annotators and GPT-4. Details are in Appendix B.4.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETUP

We evaluate 13 representative retrieval models, ranging from traditional bag-of-words models to large dense retrieval models, including the top performers from the retrieval set of the MTEB leaderboard (Muennighoff et al., 2023), BEIR (Thakur et al., 2021). First, we employ BM25 (Robertson et al., 2009) as our primary sparse, lexical-based retrieval model, which demonstrates strong performance on BEIR (Thakur et al., 2021), comparable to that of larger trained dense retrieval models. We also evaluate a diverse set of open-source dense retrieval models: the small (<1B) models are SentenceBERT (109M; Reimers & Gurevych, 2019), BGE (335M; Xiao et al., 2023), and Instructor-Large (335M Su et al., 2022), and the large (>1B) models are Instructor-XL (1.5B; Su et al., 2022), E5-Mistral (7.1B; Wang et al., 2023a), SFR-Embedding-Mistral (7.1B; Meng et al., 2024), GritLM (7.1B; Muennighoff et al., 2024), and gte-Qwen1.5 (7.7B; Li et al., 2023b). Notably, all large dense models and Instructor-Large are instruction-tuned. Lastly, we include proprietary models from Cohere (Cohere), Voyage (Voyage AI), OpenAI (OpenAI), and Google(1.2B) (Lee et al.,

<sup>9</sup>proofwiki.org; a collection of over 20K formal definitions and proofs of mathematical theorems.




Figure 3: Average nDCG@10 score on BRIGHT when using the original query vs. reasoning steps generated by GritLM, Llama3-70B and GPT-4 for retrieval. Searching with LLM reasoning steps significantly improves performance. Surprisingly, BM25 achieves the best performance in the leaderboard using reasoning steps written by GPT-4 as new queries. Detailed scores are in Table 34 to 38.

Retriever	Reranker	$k$	nDCG@10
BM25	None	-	14.3
	MiniLM	10	13.1
	MiniLM	100	8.3
	Gemini	10	15.7
	GPT-4	10	17.4
	GPT-4	100	17.0
Google	None	-	19.5
	MiniLM	10	16.0
	MiniLM	100	9.4
	Gemini	10	20.1
	GPT-4	10	21.5
	GPT-4	100	22.6

Table 3: Average reranking performance on BRIGHT. We also include the retrieval results (reranker=None) for comparison. Detailed scores can be found in Table 41 and 42.

2024). We provide details of each model in Appendix A.1. Following prior work (Thakur et al., 2021; Bajaj et al., 2018; Voorhees & Tice, 2000), we use nDCG@10 as the main metric. Please find the computing resources in Appendix A.2.

## 4.2 MAIN RESULTS

**Existing retrieval systems perform poorly on BRIGHT.** Results in Table 2 show that BRIGHT is very challenging, with the best model achieving only 24.3 nDCG@10. Although BM25 matches the < 1B models, it significantly underperforms larger models. This suggests that traditional keyword matching (“level 1 search”) is insufficient for BRIGHT. Although larger models that have been trained on semantic-based retrieval datasets like MS MARCO (Figure 1), such as GritLM (Muennighoff et al., 2024), perform better than BM25, they are still unable to solve BRIGHT. Proprietary models perform similarly to large open-source ones. Overall, the low performance indicates that the existing retrieval system cannot perform reasoning-intensive retrieval, and new methods are required to solve “level 3 search”.

**Querying with LLM reasoning steps improves retrieval performance.** Considering the strong reasoning abilities of LLMs, we propose that leveraging LLM-generated reasoning steps as queries may enhance retrieval performance. To validate this hypothesis, we prompt LLMs to write reasoning traces given a query with the following prompt: “(1) Identify the essential problem in the post. (2) Think step by step to reason about what should be included in the relevant documents. (3) Draft an answer.”. We encourage LLMs to first understand the question by summarizing it, then use chain-of-thought reasoning (Wei et al., 2022) to identify relevant content, and finally write a candidate answer. We use then use these reasoning-enhanced as new queries to evaluate all retrieval models. We use GPT-4<sup>10</sup>, GritLM (Muennighoff et al., 2024) and Llama-3-70B-Instruct<sup>11</sup> to generate reasoning steps. Figure 3 shows that using Llama-3-70B or GPT-4 reasoning steps as queries significantly improves performance compared to the original query (the detailed scores are in Tables 34 to 38). GritLM-generated reasoning steps improve BM25 performance but are less effective for other models likely due to having fewer parameters. Overall, BM25 improves the most, possibly because BM25 can adapt to different queries, while LLM-generated queries are out-of-distribution for trained models. With the best score still being below 30, significant room remains for improvement on BRIGHT.

**Retrieval augmentation boosts performance in question-answering.** An important application of retrieval is to ultimately improve the question-answering (QA) results. In addressing a StackExchange post, one typical pipeline is to first retrieve useful documents, and then answer questions leveraging

<sup>10</sup>gpt-4-0125-preview

<sup>11</sup><https://huggingface.co/meta-llama/Meta-Llama-3-70B-Instruct>



Table 4: **Question-answering results with different retrievers.** We use Claude-3.5-sonnet as the generation model and evaluate the answers with Claude-3.5-sonnet. We find that stronger retrieval typically results in better QA results, indicating the helpfulness of the annotated documents for addressing the posts in StackExchange. Performance with GPT-4 in Table 47 shows similar trends.

Retriever	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Average
None	61.3	63.8	56.9	56.1	51.7	64.1	50.0	57.7
BM25	63.4	64.7	56.2	53.2	59.3	65.7	56.3	59.8
SBERT	60.8	61.2	59.0	64.4	54.5	59.6	<b>60.8</b>	60.0
Qwen	<b>68.2</b>	<b>67.3</b>	<b>60.8</b>	<b>65.6</b>	<b>63.4</b>	<b>68.9</b>	55.7	<b>64.3</b>
Oracle	78.8	79.4	71.7	76.1	69.7	79.0	69.7	74.9

the retrieved content. While the second step could involve reasoning processes to derive an answer, we emphasize intensive reasoning in the first step to find documents. We evaluate the end-to-end QA performance of Claude-3.5-sonnet when augmented with documents retrieved by different models, and use Claude-3.5-sonnet to evaluate the answer correctness. As shown in Table 4, stronger retrievers generally result in better QA performance, with the top-performing retriever, Qwen, achieving a 6.6-point gain. However, using the oracle documents boosts performance by 17.2 points, highlighting substantial room for improvement in reasoning-intensive retrieval to enhance downstream QA results. See Appendix D for similar experiments on the TheoremQA and AoPS datasets.

Although better retrieval quality correlates with higher QA performance, the QA results may not always accurately capture retrieval performance due to the following reasons: 1) the generator model may not effectively comprehend or integrate the retrieved documents into its responses, and 2) the evaluator may fail to fully recognize the similarities and differences between two open-ended answers. Please check out Table 45 and 46 for inference and evaluation prompts used for LLMs.

## 5 ANALYSIS

### 5.1 RERANKING WITH LLMs ENHANCES RETRIEVAL PERFORMANCE

A common approach for improving retrieval results is to utilize powerful rerankers capable of performing joint computation over both the query and the documents. To this end, we investigate if performance on **BRIGHT** can be improved through reranking. We test this with a classical cross-encoder, MiniLM<sup>12</sup>, and LLMs to rerank the top  $k = \{10, 100\}$  retrieved documents. The cross-encoder is trained on the MS MARCO reranking task (Bajaj et al., 2016) and outputs a relevance score for each pair of query and document  $(Q, D_i)$ . Following Sun et al. (2023), we also rerank with LLMs by including the query and top-k documents in the prompt and asking the LLMs to order the documents based on their relevance to the query (detailed prompts can be found in Table 43). Table 3 shows that the traditional cross-encoder negatively impacts retrieval quality, with performance declining as more documents are reranked, suggesting that training rerankers on MS MARCO does not transfer well to **BRIGHT**. On the other hand, reranking by LLMs generally enhances performance. Stronger LLMs provide more significant improvements; for instance, based on BM25 retrieval results, Gemini-1.0 reranking increases the score by 1.4, and GPT-4 reranking enhances by 3.1 and continues to improve with higher  $k$ . LLMs can serve as an effective tool for reasoning-intensive retrieval, but the final results still highly depend on the underlying retrieval system.

### 5.2 ROBUSTNESS AGAINST DATA LEAKAGE FROM PRETRAINING

Many existing benchmarks suffer from inflated performance due to benchmark data contamination during large-scale pre-training (Zhou et al., 2023; Xu et al., 2024). We demonstrate that **BRIGHT** remains robust against such data leakage, even when retrieval documents are fully exposed during pre-training. To test this, we simulate a realistic scenario where language models encounter StackExchange data during internet-based training. We continue training GritLM (Muennighoff et al., 2024) on StackExchange data from our retrieval pool using both language modeling loss and contrastive

<sup>12</sup><https://huggingface.co/cross-encoder/ms-marco-MiniLM-L-12-v2>

learning on question-answer pairs (see Appendix A.3 for training details). This approach exposes the model to all StackExchange content in **BRIGHT**, while avoiding direct training on query-document mappings that require intensive reasoning. As shown in Table 5, the fine-tuned GritLM exhibits only a slight performance decrease, suggesting that conventional training procedures have limited impact on **BRIGHT** performance. These results demonstrate **BRIGHT**’s robustness to pre-training data leakage and highlight the need for novel approaches to advance reasoning-intensive retrieval.

Table 5: **BRIGHT is robust to massive pre-training.** By continuing training GritLM on Stack-Exchange data without showing the mapping between queries and documents, the model does not improve the average performance after learning the in-domain knowledge, indicating the importance of reasoning capabilities in the retrieval process.

	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Avg.
GritLM	25.0	32.8	19.0	19.9	17.3	11.6	18.0	20.5
Fine-tuned GritLM	21.1	25.5	18.8	30.7	12.7	12.1	21.9	20.4

### 5.3 LONG-CONTEXT RETRIEVAL WITH A REDUCED SEARCH SPACE IS CHALLENGING

Retrieving information from long documents is crucial for applications such as legal contracts, company financial documents, and patient notes (Saad-Falcon et al., 2024; Zhu et al., 2024). To evaluate retrieval models on reasoning-intensive tasks involving lengthy documents, we convert the StackExchange datasets to a long-context retrieval setting, where documents are complete web pages with significantly more tokens but fewer total number of documents (Table 40). With most datasets containing only a few hundred documents, nDCG@10, which evaluates the top 10 results, becomes more susceptible to randomness. Moreover, processing 10 long documents with an average length of up to 40,000 tokens is challenging for both humans and LLMs. Therefore, we decide to use recall@1 metric to provide a more reliable measure in this setting. Table 6 presents the average scores for 8 datasets from StackExchange and Pony. The highest recall achieved is 27.8, indicating that even with significantly reduced retrieval pools, the combination of long-context documents and intensive reasoning remains challenging for existing retrieval models.

Table 6: **Long-context retrieval performance where retrievers retrieve from unsplit web pages.** The results are reported as the average recall@1 score of StackExchange and Pony datasets. More detailed numbers can be found in Table 39.

BM25	BGE	Inst-L	SBERT	E5	SFR	Inst-XL	GritLM	Qwen	Cohere	OpenAI	Voyage	Google
11.4	14.8	18.2	17.4	25.5	26.0	17.8	26.0	27.8	18.4	21.9	24.6	22.4

## 6 CONCLUSION

We introduce **BRIGHT**, the first retrieval benchmark that encompasses realistic retrieval scenarios requiring intensive reasoning steps to identify relevant documents. We utilize existing online document structures and dedicate substantial human effort to curate **BRIGHT** and verify its correctness. Through extensive evaluation, we find that existing retrieval models perform extremely poorly on **BRIGHT**, with a maximum nDCG@10 score of only 24.3. Augmenting retrieval queries with reasoning steps generated by LLMs improves performance, but even the best model still achieves a score below 30. Furthermore, strong retrieval results can significantly improve downstream performance on reasoning tasks, highlighting a practical application of reasoning-intensive retrieval. We hope that **BRIGHT** will contribute to future research investigations into pushing the state-of-the-art in this direction.

## 7 CODE OF ETHICS AND ETHICS STATEMENT

In the process of collecting datasets for BRIGHT, we ensure that all sources come from public data, used solely for academic research and not for commercial purposes, in full compliance with the copyright rights granted by the sources. We guarantee that none of the datasets contain harmful information to society, such as racial discrimination, violence, or any private data. Our work aims to contribute to the welfare of society and humanity, and any researcher is free to use our dataset for research purposes. All the data and experiments presented in our paper adhere to the highest standards of scientific excellence, ensuring the authenticity and accuracy of the data.

## 8 REPRODUCIBILITY

Our datasets and annotation process are introduced in Sec. 3, and the experimental settings are described in Sec. 4. Specific implementation details can be found in App. A. To facilitate the reproduction of our experiments, the code and data are provided in <https://drive.google.com/file/d/1repp1sMj17pfV-nrivpvb4tJg4pj0sm/view?usp=sharing>.

## REFERENCES

- Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hanna Hajishirzi, and Wen tau Yih. Reliable, adaptable, and attributable language models with retrieval. 2024. url <https://api.semanticscholar.org/corpusid:268248911>. opra, and colin raffel. extracting training data from large language models. In *USENIX Security Symposium (USENIX)*, 2020.
- Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, and Wen-tau Yih. Task-aware retrieval with instructions. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 3650–3675, Toronto, Canada, July 2023a. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.225. URL <https://aclanthology.org/2023.findings-acl.225>.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*, 2023b.
- Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi, and Wen tau Yih. Reliable, adaptable, and attributable language models with retrieval, 2024.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. Ms marco: A human generated machine reading comprehension dataset, 2018.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack Rae, Erich Elsen, and Laurent Sifre. Improving language models by retrieving from trillions of tokens. In *International Conference on Machine Learning (ICML)*, volume 162, pp. 2206–2240, 2022. URL <https://proceedings.mlr.press/v162/borgeaud22a.html>.

- 594 Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading Wikipedia to answer  
595 open-domain questions. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th*  
596 *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.  
597 1870–1879, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi:  
598 10.18653/v1/P17-1171. URL <https://aclanthology.org/P17-1171>.
- 599 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared  
600 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large  
601 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- 602  
603 Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and  
604 Tony Xia. TheoremQA: A theorem-driven question answering dataset. In Houda Bouamor, Juan  
605 Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural*  
606 *Language Processing*, pp. 7889–7901, Singapore, December 2023a. Association for Computational  
607 Linguistics. doi: 10.18653/v1/2023.emnlp-main.489. URL [https://aclanthology.org/](https://aclanthology.org/2023.emnlp-main.489)  
608 [2023.emnlp-main.489](https://aclanthology.org/2023.emnlp-main.489).
- 609 Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and  
610 Tony Xia. Theoremqa: A theorem-driven question answering dataset. In *The 2023 Conference on*  
611 *Empirical Methods in Natural Language Processing*, 2023b.
- 612  
613 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
614 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
615 Schulman. Training verifiers to solve math word problems, 2021.
- 616  
617 Cohere. Cohere-embed-english-light-v3.0, 2022. URL [https://huggingface.co/Cohere/](https://huggingface.co/Cohere/Cohere-embed-english-light-v3.0)  
618 [Cohere-embed-english-light-v3.0](https://huggingface.co/Cohere/Cohere-embed-english-light-v3.0).
- 619 Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. Overview of  
620 the trec 2019 deep learning track, 2020.
- 621  
622 Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In  
623 *International Conference on Learning Representations (ICLR)*, 2024.
- 624  
625 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and  
626 memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing*  
627 *Systems (NeurIPS)*, 2022.
- 628 Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus  
629 Leippold. Climate-fever: A dataset for verification of real-world climate claims. *arXiv preprint*  
630 *arXiv:2012.00614*, 2020.
- 631  
632 Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. Eli5:  
633 Long form question answering. *arXiv preprint arXiv:1907.09190*, 2019.
- 634  
635 Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. Enabling large language models to generate  
636 text with citations. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- 637  
638 Faegheh Hasibi, Fedor Nikolaev, Chenyan Xiong, Krisztian Balog, Svein Erik Bratsberg, Alexander  
639 Kotov, and Jamie Callan. Dbpedia-entity v2: a test collection for entity search. In *Proceedings*  
640 *of the 40th International ACM SIGIR Conference on Research and Development in Information*  
*Retrieval*, pp. 1265–1268, 2017.
- 641  
642 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
643 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In  
644 *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*  
645 *(Round 2)*, 2021. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.
- 646  
647 Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt.  
CodeSearchNet challenge: Evaluating the state of semantic code search. *arXiv preprint*  
*arXiv:1909.09436*, 2019.

- 648 Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand  
649 Joulin, and Edouard Grave. Unsupervised dense information retrieval with contrastive learn-  
650 ing. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=jKN1pXi7b0>.  
651  
652
- 653 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando  
654 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free  
655 evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
- 656 Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi  
657 Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. *arXiv*  
658 *preprint arXiv:2004.04906*, 2020.  
659
- 660 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris  
661 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a  
662 benchmark for question answering research. *Transactions of the Association for Computational*  
663 *Linguistics*, 7:453–466, 2019.
- 664 Jinhyuk Lee, Zhuyun Dai, Xiaoqi Ren, Blair Chen, Daniel Cer, Jeremy R Cole, Kai Hui, Michael  
665 Boratko, Rajvi Kapadia, Wen Ding, et al. Gecko: Versatile text embeddings distilled from large  
666 language models. *arXiv preprint arXiv:2403.20327*, 2024.  
667
- 668 Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. Latent retrieval for weakly supervised open  
669 domain question answering. In *Association for Computational Linguistics (ACL)*, pp. 6086–6096,  
670 2019. doi: 10.18653/v1/P19-1612. URL <https://aclanthology.org/P19-1612>.
- 671 Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem.  
672 Camel: Communicative agents for "mind" exploration of large language model society. In *Thirty-*  
673 *seventh Conference on Neural Information Processing Systems*, 2023a.  
674
- 675 Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards  
676 general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*,  
677 2023b.
- 678 Zhuoqun Li, Hongyu Lin, Tianshu Wang, Boxi Cao, Yaojie Lu, Weixiang Zhou, Hao Wang, Zhenyu  
679 Zeng, Le Sun, and Xianpei Han. Url: Universal referential knowledge linking via task-instructed  
680 representation compression. *arXiv preprint arXiv:2404.16248*, 2024.  
681
- 682 Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by ratio-  
683 nale generation: Learning to solve and explain algebraic word problems. In Regina Barzilay  
684 and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association*  
685 *for Computational Linguistics (Volume 1: Long Papers)*, pp. 158–167, Vancouver, Canada,  
686 July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1015. URL  
687 <https://aclanthology.org/P17-1015>.
- 688 Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk,  
689 and Alexandra Balahur. Www’18 open challenge: financial opinion mining and question answering.  
690 In *Companion proceedings of the the web conference 2018*, pp. 1941–1942, 2018.  
691
- 692 Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming Xiong, Yingbo Zhou, and Semih Yavuz.  
693 Sfrembedding-mistral: enhance text retrieval with transfer learning. *Salesforce AI Research*  
694 *Blog*, 3, 2024.
- 695 Niklas Muennighoff. Sgpt: Gpt sentence embeddings for semantic search. *arXiv preprint*  
696 *arXiv:2202.08904*, 2022.  
697
- 698 Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. MTEB: Massive text embed-  
699 ding benchmark. In Andreas Vlachos and Isabelle Augenstein (eds.), *Proceedings of the 17th*  
700 *Conference of the European Chapter of the Association for Computational Linguistics*, pp. 2014–  
701 2037, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/  
v1/2023.eacl-main.148. URL <https://aclanthology.org/2023.eacl-main.148>.



- 702 Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and  
703 Douwe Kiela. Generative representational instruction tuning. *arXiv preprint arXiv:2402.09906*,  
704 2024.
- 705  
706 Hanseok Oh, Hyunji Lee, Seonghyeon Ye, Haebin Shin, Hansol Jang, Changwook Jun, and Minjoon  
707 Seo. Instructir: A benchmark for instruction following of information retrieval models, 2024.
- 708  
709 OpenAI. New embedding models and api updates, 2024. URL [https://openai.com/index/  
710 new-embedding-models-and-api-updates/](https://openai.com/index/new-embedding-models-and-api-updates/).
- 711  
712 Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao,  
713 James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim  
714 Rocktäschel, and Sebastian Riedel. KILT: a benchmark for knowledge intensive language tasks.  
715 In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven  
716 Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021  
717 Conference of the North American Chapter of the Association for Computational Linguistics:  
718 Human Language Technologies*, pp. 2523–2544, Online, June 2021. Association for Computational  
719 Linguistics. doi: 10.18653/v1/2021.naacl-main.200. URL [https://aclanthology.org/  
2021.naacl-main.200](https://aclanthology.org/2021.naacl-main.200).
- 720  
721 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language  
722 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 723  
724 Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks.  
725 *arXiv preprint arXiv:1908.10084*, 2019.
- 726  
727 Stephen Robertson, Hugo Zaragoza, et al. The probabilistic relevance framework: Bm25 and beyond.  
728 *Foundations and Trends® in Information Retrieval*, 3(4):333–389, 2009.
- 729  
730 Jon Saad-Falcon, Daniel Y Fu, Simran Arora, Neel Guha, and Christopher Ré. Benchmarking and  
731 building long-context retrieval models with loco and m2-bert. *arXiv preprint arXiv:2402.07440*,  
732 2024.
- 733  
734 Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettle-  
735 moyer, and Wen tau Yih. REPLUG: Retrieval-augmented black-box language models, 2024.
- 736  
737 Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih,  
738 Noah A Smith, Luke Zettlemoyer, and Tao Yu. One embedder, any task: Instruction-finetuned text  
739 embeddings. *arXiv preprint arXiv:2212.09741*, 2022.
- 740  
741 Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih,  
742 Noah A. Smith, Luke Zettlemoyer, and Tao Yu. One embedder, any task: Instruction-finetuned  
743 text embeddings. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of  
744 the Association for Computational Linguistics: ACL 2023*, pp. 1102–1121, Toronto, Canada, July  
745 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.71. URL  
746 <https://aclanthology.org/2023.findings-acl.71>.
- 747  
748 Hongjin Su, Shuyang Jiang, Yuhang Lai, Haoyuan Wu, Boao Shi, Che Liu, Qian Liu, and Tao Yu.  
749 Arks: Active retrieval in knowledge soup for code generation. *arXiv preprint arXiv:2402.12317*,  
750 2024.
- 751  
752 Weiwei Sun, Lingyong Yan, Xinyu Ma, Pengjie Ren, Dawei Yin, and Zhaochun Ren. Is chat-  
753 gpt good at search? investigating large language models as re-ranking agent. *arXiv preprint  
754 arXiv:2304.09542*, 2023.
- 755  
756 Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. Beir: A  
757 heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint  
758 arXiv:2104.08663*, 2021.
- 759  
760 James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale  
761 dataset for fact extraction and verification. *arXiv preprint arXiv:1803.05355*, 2018.

- 756 Ellen M. Voorhees and Dawn M. Tice. Building a question answering test collection. In *Proceedings*  
757 *of the 23rd Annual International ACM SIGIR Conference on Research and Development in*  
758 *Information Retrieval, SIGIR '00*, pp. 200–207, New York, NY, USA, 2000. Association for  
759 Computing Machinery. ISBN 1581132263. doi: 10.1145/345508.345577. URL <https://doi.org/10.1145/345508.345577>.
- 761 Voyage AI. Embeddings, 2024. URL <https://docs.voyageai.com/docs/embeddings>.
- 762 David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and  
763 Hannaneh Hajishirzi. Fact or fiction: Verifying scientific claims. *arXiv preprint arXiv:2004.14974*,  
764 2020.
- 765 Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. Improving  
766 text embeddings with large language models. *arXiv preprint arXiv:2401.00368*, 2023a.
- 767 Xiaoyue Wang, Jianyou Wang, Weili Cao, Kaicheng Wang, Ramamohan Paturi, and Leon Bergen.  
768 Birco: A benchmark of information retrieval tasks with complex objectives. *arXiv preprint*  
769 *arXiv:2402.14151*, 2024.
- 770 Zengzhi Wang, Rui Xia, and Pengfei Liu. Generative ai for math: Part i – mathpile: A billion-token-  
771 scale pretraining corpus for math, 2023b. URL <https://arxiv.org/abs/2312.17120>.
- 772 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
773 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
774 *neural information processing systems*, 35:24824–24837, 2022.
- 775 Orion Weller, Benjamin Chang, Sean MacAvaney, Kyle Lo, Arman Cohan, Benjamin Van Durme,  
776 Dawn Lawrie, and Luca Soldaini. Followir: Evaluating and teaching information retrieval models  
777 to follow instructions. *arXiv preprint arXiv:2403.15246*, 2024.
- 778 Chenghao Xiao, G Thomas Hudson, and Noura Al Moubayed. Rar-b: Reasoning as retrieval  
779 benchmark, 2024.
- 780 Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to  
781 advance general chinese embedding, 2023.
- 782 Ruijie Xu, Zengzhi Wang, Run-Ze Fan, and Pengfei Liu. Benchmarking benchmark leakage in large  
783 language models. *arXiv preprint arXiv:2404.18824*, 2024.
- 784 Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil,  
785 Ryan Prenger, and Anima Anandkumar. Leandojo: Theorem proving with retrieval-augmented  
786 language models, 2023.
- 787 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov,  
788 and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question  
789 answering. *arXiv preprint arXiv:1809.09600*, 2018.
- 790 Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou,  
791 and Jingren Zhou. Scaling relationship on learning mathematical reasoning with large language  
792 models, 2023.
- 793 Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhao Chen.  
794 Mammoth: Building math generalist models through hybrid instruction tuning, 2023.
- 795 Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao, Pranav  
796 Raja, Dylan Slack, Qin Lyu, et al. A careful examination of large language model performance on  
797 grade school arithmetic. *arXiv preprint arXiv:2405.00332*, 2024.
- 798 Kun Zhou, Yutao Zhu, Zhipeng Chen, Wentong Chen, Wayne Xin Zhao, Xu Chen, Yankai Lin,  
799 Ji-Rong Wen, and Jiawei Han. Don’t make your llm an evaluation benchmark cheater. *arXiv*  
800 *preprint arXiv:2311.01964*, 2023.
- 801 Dawei Zhu, Liang Wang, Nan Yang, Yifan Song, Wenhao Wu, Furu Wei, and Sujian Li. Longembed:  
802 Extending embedding models for long context retrieval, 2024.

810	CONTENTS	
811		
812		
813	<b>1 Introduction</b>	<b>1</b>
814		
815	<b>2 Related Work</b>	<b>2</b>
816		
817	<b>3 Constructing BRIGHT</b>	<b>3</b>
818		
819	3.1 Task formulation . . . . .	3
820	3.2 StackExchange: Retrieving web pages that help answer questions . . . . .	4
821		
822	3.3 Coding: Retrieving documentation or similar solved problems . . . . .	5
823		
824	3.4 Theorem-based questions: Retrieve solved problems using the same theorems or relevant theorem statements . . . . .	6
825		
826		
827	<b>4 Experiments</b>	<b>7</b>
828	4.1 Experimental setup . . . . .	7
829		
830	4.2 Main results . . . . .	8
831		
832	<b>5 Analysis</b>	<b>9</b>
833		
834	5.1 Reranking with LLMs enhances retrieval performance . . . . .	9
835	5.2 Robustness against data leakage from pretraining . . . . .	9
836	5.3 Long-context retrieval with a reduced search space is challenging . . . . .	10
837		
838		
839	<b>6 Conclusion</b>	<b>10</b>
840		
841	<b>7 Code of Ethics and Ethics statement</b>	<b>11</b>
842		
843	<b>8 Reproducibility</b>	<b>11</b>
844		
845		
846	<b>A Experiment Details</b>	<b>17</b>
847		
848	A.1 Models and Instructions . . . . .	17
849	A.2 Computing Resources . . . . .	18
850	A.3 Continual Training Setup . . . . .	19
851		
852		
853	<b>B Dataset Construction</b>	<b>19</b>
854	B.1 StackExchange . . . . .	19
855	B.2 STEM question and solution corpus for TheoremQA and AoPS . . . . .	20
856		
857	B.3 TheoremQA: Rephrasing questions into specific scenarios . . . . .	20
858	B.4 TheoremQA: Annotating relevant theorems . . . . .	20
859		
860	B.5 AoPS: Connecting AoPS problems to the MATH dataset . . . . .	21
861	B.6 LeetCode . . . . .	21
862		
863		
	<b>C Data Examples</b>	<b>22</b>

864	<b>D Downstream performance</b>	<b>22</b>
865		
866	<b>E Annotator instructions</b>	<b>22</b>
867		
868	E.1 StackExchange . . . . .	22
869	E.2 TheoremQA . . . . .	23
870		
871	E.3 AoPS . . . . .	23
872	E.4 LeetCode . . . . .	23
873		
874	E.5 LLM usage . . . . .	24
875	E.6 Sensitive information . . . . .	24
876		
877	<b>F Limitations and Future Work</b>	<b>24</b>
878		
879	<b>G Potential social impacts</b>	<b>24</b>
880		
881	<b>H Comparison to RAR-b</b>	<b>24</b>
882		

## A EXPERIMENT DETAILS

### A.1 MODELS AND INSTRUCTIONS

For each model used in this paper, Table 7 provides information on the size, architecture, maximum context length of queries and documents, whether we include instructions and the specific version we use in the experiments. All parameters are set by following the official tutorial. The only exceptions are Inst-L and Inst-XL, where we empirically find that extending the maximum context length to 2048 significantly enhances the performance. In Table 8, 9, 10 and 11, We specify the instructions used for BGE, Inst-L, Inst-XL, E5, GritLM, Qwen and SFR in each dataset. For the embedding model from Google, we use the parameter "task" with the values "RETRIEVAL\_QUERY" and "RETRIEVAL\_DOCUMENT" to distinguish queries from documents and use the parameter "input\_type" with the values "query" and "document" for the embedding model from Voyage.

Table 7: **All 13 models benchmarked in experiments.** We report the number of parameters of each model except the sparse model BM25 and proprietary models without public information. Regarding the model architecture, we distinguish between sparse and dense models and further classify dense models as encoders or decoders if known. Max  $|Q|$  and Max  $|D|$  denote the maximum context length we use for each model in the experiments. The instruction column indicates whether we include instructions in the retrieval. The version column denotes the specific checkpoint or implementation.

	Size	Architecture	Max $ Q $	Max $ D $	Instruction	Version	License
<i>Sparse model</i>							
BM25 (Robertson et al., 2009)	N/A	Sparse	$\infty$	$\infty$	No	gensim <sup>13</sup>	LGPL-2.1-only
<i>Open-sourced models (&lt;1B)</i>							
SBERT (Reimers & Gurevych, 2019)	109M	Encoder	512	512	No	all-mpnet-base-v2	Apache-2.0
BGE (Xiao et al., 2023)	335M	Encoder	512	512	No	bge-large-en-v1.5	MIT
Inst-L Su et al. (2022)	335M	Encoder	2048	2048	Yes	instructor-large	Apache-2.0
<i>Open-sourced models (&gt;1B)</i>							
Inst-XL (Su et al., 2022)	1.5B	Encoder	2048	2048	Yes	instructor-xl	Apache-2.0
E5 (Wang et al., 2023a)	7.1B	Decoder	4096	4096	Yes	e5-mistral-7b-instruct	MIT
GritLM (Muennighoff et al., 2024)	7.1B	Decoder	256	2048	Yes	GritLM-7B	Apache-2.0
SFR (Meng et al., 2024)	7.1B	Decoder	4096	4096	Yes	SFR-Embedding-Mistral	CC-BY-NC-4.0
Qwen (Li et al., 2023b)	7.7B	Decoder	8192	8192	Yes	gte-Qwen1.5-7B-instruct	Apache-2.0
<i>Proprietary models</i>							
Cohere (Cohere)	N/A	Dense	512	512	No	Cohere-embed-english-v3.0	Company
Google (Lee et al., 2024)	1.2B	Dense	2000	2000	Yes	text-embedding-preview-0409, dimension=768	Company
OpenAI (OpenAI)	N/A	Dense	8191	8191	No	text-embedding-3-large	Company
Voyage (Voyage AI)	N/A	Dense	16000	16000	Yes	voyage-large-2-instruct	Company

918 Table 8: **Instructions used for benchmarking StackExchange datasets.** {domain} is one of  
 919 Biology, Earth Science, Economics, Psychology, Robotics, Stack Overflow and Sustainable Living.

920 Models	921 Instructions
922 BGE	<b>Query:</b> Represent this {domain} post for searching relevant passages:
923 Inst-L, 924 Inst-XL	<b>Query:</b> Represent the {domain} post for retrieving relevant paragraphs: <b>Doc:</b> Represent the {domain} paragraph for retrieval:
925 E5, GritLM, 926 Qwen, SFR	<b>Query:</b> Given a {domain} post, retrieve relevant passages that help answer the post

928  
929  
930 Table 9: **Instructions used for benchmarking the LeetCode dataset.**

931 Models	932 Instructions
933 BGE	<b>Query:</b> Represent this Coding problem for searching relevant examples:
934 Inst-L, 935 Inst-XL	<b>Query:</b> Represent the Coding problem for retrieving relevant examples: <b>Doc:</b> Represent the Coding example for retrieval:
936 E5, GritLM, 937 Qwen, SFR	<b>Query:</b> Given a Coding problem, retrieve relevant examples that help answer the 938 problem

939  
940 Table 10: **Instructions used for benchmarking the Pony dataset.**

941 Models	942 Instructions
943 BGE	<b>Query:</b> Represent this Pony question for searching relevant passages:
944 Inst-L, 945 Inst-XL	<b>Query:</b> Represent the Pony question for retrieving relevant paragraphs: <b>Doc:</b> Represent the Pony paragraph for retrieval:
946 E5, GritLM, 947 Qwen, SFR	<b>Query:</b> Given a Pony question, retrieve relevant passages that help answer the 948 question

949  
950 Table 11: **Instructions used for benchmarking Math datasets (AoPS and TheoremQA).**

951 Models	952 Instructions
953 BGE	<b>Query:</b> Represent this Math problem for searching relevant examples:
954 Inst-L, 955 Inst-XL	<b>Query:</b> Represent the Math problem for retrieving relevant examples: <b>Doc:</b> Represent the Math example for retrieval:
956 E5, GritLM, 957 Qwen, SFR	<b>Query:</b> Given a Math problem, retrieve relevant examples that help answer the 958 problem

## 961 A.2 COMPUTING RESOURCES

962  
963 We run all experiments on NVIDIA V100, A100, or H100 GPUs. The amount of time that it takes  
 964 to complete one round of experiments is dependent on the model. For the sparse model, BM25,  
 965 the evaluation takes less than 1 hour on CPU-only machines. For the open-sourced dense models  
 966 ( $< 1B$ ), the evaluation requires about 8 hours on one H100 GPU. For the open-sourced dense models  
 967 ( $> 1B$ ), the evaluation takes up to 36 hours on one H100 GPU. We leverage FlashAttention (Dao  
 968 et al., 2022; Dao, 2024) for speedup when evaluating the dense models. For the proprietary models,  
 969 the evaluation speed is dependent on the API bandwidth, but we found that one round of experiments  
 970 can be completed within 2 days.  
 971



### A.3 CONTINUAL TRAINING SETUP

In Section 5.2, we introduce the continual training method GritLM on StackExchange data to evaluate whether training on in-domain data enhances the performance of **BRIGHT**. Detailed experimental settings are described in this section. Specifically, we follow GritLM to train models with two distinct objectives: a contrastive loss to maintain the model’s retrieval capability and a language modeling loss to preserve the model’s language generation ability. For training with the contrastive loss, we collect 3,200 (post, answer) pairs from the Biology, Earth Science, Economics, Psychology, Robotics, and Stack Overflow sections of StackExchange, and 1,538 pairs from Sustainable Living. Each post’s answer is used as a positive example, with other answers serving as in-batch negatives. For training with the language modeling loss, we use both positive and negative documents from each domain within the StackExchange subsection of **BRIGHT**. These documents are split into chunks of 2048 tokens, and we sample up to 3,200 chunks for training. We use a small batch size of 64 to ensure sufficient learning steps, while following the other hyperparameters as outlined in Muennighoff et al. (2024). We continue training GritLM for 10 epochs, benchmarking the checkpoint from each epoch on the StackExchange datasets of **BRIGHT**. The detailed scores are in Table 12, where we copy the scores of GritLM to epoch=0 for easier reference. The results indicate no significant improvement across the 10 epochs, suggesting that even with intensive inclusion of StackExchange data or relevant domain knowledge in the training data of language models or retrievers, performance may not increase substantially without enhancing incorporating reasoning into the retrieval process.

Table 12: **Scores of finetuned GritLM of every epoch on StackExchange datasets of BRIGHT.** Epoch=0 indicates the performance of GritLM without further training.

Epoch	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Avg.
0 (GritLM)	25.0	32.8	19.0	19.9	17.3	11.6	18.0	20.5
1	22.2	25.4	17.6	28.1	11.1	9.8	19.6	19.1
2	18.7	23.8	13.5	19.3	10.7	10.2	16.5	16.1
3	20.9	23.6	16.9	25.2	11.1	8.5	16.6	17.5
4	24.3	28.0	18.3	26.9	13.4	13.3	20.0	20.6
5	23.1	28.5	18.4	26.1	14.6	11.7	21.6	20.6
6	19.9	26.4	16.0	27.9	9.6	9.3	19.3	18.3
7	24.3	25.4	16.5	28.1	11.0	9.8	17.0	18.9
8	21.6	28.7	19.2	28.7	11.1	11.8	22.4	20.5
9	21.3	29.0	20.0	28.7	11.4	14.3	22.0	21.0
10	21.1	25.5	18.8	30.7	12.7	12.1	21.9	20.4

## B DATASET CONSTRUCTION

### B.1 STACKEXCHANGE

**Passage split.** To split web pages or long documents into smaller pieces of passages, we employ simple heuristics with separators like two new line symbols, "#" in markdown files without additional assumptions on file structures. Although this split may not be optimal for every document, it simulates the realistic setting where long documents are automatically processed without human or expert intervention.

**False positive and false negative.** We discuss the rationale for avoiding false positives and false negatives in the data collection process. In StackExchange, the positive relevance between queries and documents is manually verified by annotators, with detailed reasoning traces to illustrate their thinking process. To avoid false negatives, the annotators only select a StackExchange post that clearly distinguishes from previously annotated examples in the same domain, e.g., different entities and semantic meanings in both posts and answers, etc. This ensures that no pair of examples share the same positive documents.

<sup>13</sup>Specifically, we use the LuceneBM25Model from `gensim` and the text analyzer from `pyserini`

## B.2 STEM QUESTION AND SOLUTION CORPUS FOR THEOREMQA AND AOPS

In this subsection, we describe the construction of the STEM question and solution corpus, which is used for both TheoremQA Questions and AoPS. We source the documents (pairs of problem statements and solutions)  $D_i = (Q_i, A_i)$  from existing datasets—GSM8K (Cobbe et al., 2021), GSM8K-RFT (Yuan et al., 2023), MATH (Hendrycks et al., 2021), AQuA-RAT (Ling et al., 2017), TheoremQA (Chen et al., 2023a), and CAMEL-Math (Li et al., 2023a). To reduce the likelihood of false negatives among the STEM corpus, we leverage the metadata from the original datasets to exclude specific documents from the corpus for each test query. For example, CAMEL-Math contains problem-solution pairs labeled with the category “Calculus”, which covers different questions that involve derivatives and integrals. Therefore, for queries in TheoremQA that uses “derivative chain rule” or “integral rules”, we excluded CAMEL-Math pairs in the category “Calculus” to reduce possible false negatives. Thus, for each test query in TheoremQA-Q and AoPS, we manually decide which labels in the other datasets to exclude based on the metadata. We do not exclude any problem-solution pairs from GS8K, GSM8K-RFT, or AQuA-RAT due to the relative elementary difficulty (mostly basic algebra questions) in comparison to our test queries, which leverages more advanced theorems and techniques. The mapping from the test query category to the excluded problem-solution categories can be found in Table 14, 15, and 16.

An alternative approach to excluding false negatives from the corpus for each test query is to inspect every problem-solution pair and annotate if they are relevant to test query. Although this would yield harder negatives and additional positives, we opt to not to use this approach due to its expensive cost to conduct annotation between every test query and possible candidates.

## B.3 THEOREMQA: REPHRASING QUESTIONS INTO SPECIFIC SCENARIOS

TheoremQA is a dataset consisting of theorem-driven questions in mathematics, physics, finance, and computer science and electrical engineering (Chen et al., 2023a). For each question in TheoremQA, we refer it to MathInstruct dataset<sup>14</sup> (Yue et al., 2023), as each question in this dataset is annotated with the reasoning steps and final answers.

From preliminary analysis, we found that TheoremQA questions are often written in a way such that the theorem used to solve the problem is explicitly mentioned in the question. As a result, questions that use the same theorem can have high keyword overlap, which means retrievers can easily retrieve the correct document by matching the keywords. Thus, we rewrite the questions in TheoremQA by grounding them in real-world scenarios or applications, which makes the reasoning steps less explicit and provides more diverse questions. We leverage GPT-4 with manually written instructions and in-context demonstrations to rewrite the queries  $Q$ . We provide the prompt used to rewrite TheoremQA questions and an example in Table 13. After rewriting the question, the authors manually inspect each rewritten question to ensure that the question is solvable and consistent with the original question (i.e., the reasoning steps and final answer still hold). When applicable, we manually edit the rewritten question to improve the fluency and coherence of the question, and discard the query if the rewritten question is not solvable or consistent with the original question. Consequently, we obtain 206 rewritten questions in TheoremQA from the original set of 800 questions.

## B.4 THEOREMQA: ANNOTATING RELEVANT THEOREMS

For TheoremQA-Theorems, we use the test queries from TheoremQA-Questions and annotate them with useful theorem proofs and definitions. We source mathematical theorem proofs and definitions from ProofWiki, which is community-driven effort with more than 20K formal definitions and proofs of mathematical theorems. ProofWiki is preprocessed and provided by MathPile (Wang et al., 2023b). We opt to map the original theorem names to the documents in ProofWiki so that the gold documents have consistent forms with other documents in the corpus.

For each test query, we first construct a candidate set of useful documents from ProofWiki using the theorem name and definition provided by the original TheoremQA dataset. Specifically, we construct the candidate set with the following steps:

<sup>14</sup><https://huggingface.co/datasets/TIGER-Lab/MathInstruct>

- 1080 1. Find documents where the theorem name exists as a substring. We discard this set if there  
1081 are more than 100 such documents, which typically means that the theorem name is too  
1082 common.
- 1083 2. Using the theorem name and definition from the original dataset as the query, we use BM25  
1084 to retrieve the top  $k = 10$  documents from ProofWiki.  
1085

1086 Then, we prompt GPT-4 (gpt-4-0125-preview) to check if each document’s described theorem  
1087 are used in the problem solutions, which labels each candidate as either a positive or negative  
1088 document. The prompt for this step can be found in Table 17. The authors manually annotated 50  
1089 instances and found a substantial agreement of Cohen’s  $\kappa = 0.62$  with the model judgments. Finally,  
1090 we keep test queries with at least one positive document.

## 1091 B.5 AOPS: CONNECTING AOPS PROBLEMS TO THE MATH DATASET

1093 AoPS Wiki is a community-driven platform where users can post problems and solutions to math  
1094 competition problems. These math competitions include, but are not limited to the American  
1095 Mathematics Competitions (AMC), the American Invitational Mathematics Examination (AIME),  
1096 and the International Mathematical Olympiad (IMO). In addition to problems, the AoPS Wiki also  
1097 contains articles on various topics in mathematics, such as Fermat’s Little Theorem and Ball and  
1098 Urns. These articles not only describe the theorem or technique but also link to problems that can  
1099 be solved by them. We browse the AoPS Wiki and collect the topics and the linked problems. The  
1100 topics are listed in Table 18.

1101 Although math competition problems are used in previous datasets, such as MATH (Hendrycks  
1102 et al., 2021), they lack the necessary annotations on the problem-solving skills to construct positive  
1103 documents. Thus, we opt to collect these annotations from AoPS Wiki instead.

1104 Furthermore, since MATH examples are used in the STEM corpus, we deduplicate them by matching  
1105 the collected problems with the MATH instances. Specifically, for each question  $Q$  we collected from  
1106 AoPS, we find the closest problem statement in MATH using n-gram overlap, and manually check  
1107 if they are the same problem. If the same problems are found, we merge them into one instance,  
1108 otherwise, we create a new instance and insert it into the corpus.

## 1110 B.6 LEETCODE

1111 We first obtain the publicly available LeetCode<sup>15</sup> questions from HuggingFace<sup>16</sup>. Our retrieval pool  
1112 is sourced from a combination of LeetCode and CodeSearchNet, including a problem description  
1113 and a solution in each example. In the following sections, we outline the process for constructing  
1114 positive examples and performing additional checks to minimize the likelihood of false negatives  
1115 while ensuring the use of a large retrieval pool.  
1116

1117 **Using similar questions as positive examples.** For each question, we obtain the gold pair annotations  
1118 from the “Similar Questions” field, which contains links to other LeetCode questions that are similar  
1119 to the problem. While the website does not explicitly describe the guidelines behind how this field is  
1120 populated, our qualitative analysis showed that these questions have a high overlap in terms of the  
1121 data structure, algorithms, and/or logical reasoning used to solve the problem.

### 1122 **Select problems based on real-world scenarios to avoid false negatives from CodeSearchNet.**

1123 From a preliminary qualitative analysis, we discovered that some questions in LeetCode are frequently  
1124 found in CodeSearchNet due to the popularity of certain algorithms and the simplicity of their problem  
1125 statements. Examples include implementing sorting algorithms or merging two linked lists, which  
1126 could unexpectedly introduce false positives into this retrieval setup. Thus, we use an additional  
1127 filtering step—we only keep questions that are grounded in real-world concepts that are not as  
1128 commonly used in the context of coding problems. The intuition behind this is similar to the  
1129 TheoremQA annotations process: the reasoning steps (i.e., the algorithms and data structures used to  
1130 solve the problem) cannot be as easily deciphered as if the problem statement clearly describes the  
1131 algorithm and data structure used. To this end, we first manually write instructions with six-shot in-  
1132 context learning demonstrations, and use GPT-4 (gpt-4-0125-preview) to classify all LeetCode

1133 <sup>15</sup><https://leetcode.com/>

<sup>16</sup><https://huggingface.co/datasets/greengranger/leetcode>

1134 questions. Then, we validate the GPT-4 judgment with 80 annotations by the authors. Authors’  
 1135 and GPT-4’s annotations have a Cohen’s kappa of 0.73, which suggests substantial agreement. The  
 1136 prompt used in this step and examples can be found in Table 19. GPT-4 judged a total of 291 samples  
 1137 of being grounded in real-world concepts, and we randomly sample 142 questions from this set to  
 1138 construct our test set.

### 1139 **Remove examples with high topic overlap from the pool to avoid false negatives in LeetCode.**

1140 To avoid potential false negatives that are not annotated by the LeetCode website, we leverage the  
 1141 “Topics” field from the website, which contains information about the algorithms and data structure  
 1142 used in the problem, such as “stack” and “breadth first search”. For each question  $Q$ , we collect its  
 1143 topics  $T(Q) = \{t_q^1, \dots, t_q^m\}$  from the LeetCode website, where  $t_q$  denotes the  $m \geq 1$  different topics  
 1144 assigned to the question by the website. Since each question may have multiple tags, we exclude  
 1145 other questions that have a high overlap with test questions from the corpus for that specific question.  
 1146 Specifically, we exclude question  $Q'$  from the corpus for test query  $Q$  if  $\frac{|T(Q) \cap T(Q')|}{|T(Q)|} \geq 0.5$ , because  
 1147 more than half of the topics used in  $Q'$  are also used in  $Q$ . This means that the two questions can be  
 1148 highly related in reasoning steps. Overlap smaller than the 0.5 threshold means that  $Q'$  is unlikely  
 1149 to be related, and thus a false negative, to the test question  $Q$ . Finally, we construct the rest of the  
 1150 corpus from CodeSearchNet (Husain et al., 2019)<sup>17</sup>. We only consider the Python functions as the  
 1151 solutions to the LeetCode questions are all in Python.

## 1152 C DATA EXAMPLES

1153  
 1154 In Table 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, and 31, we show more examples in **BRIGHT**.

## 1155 D DOWNSTREAM PERFORMANCE

1156  
 1157 In this work, we evaluated the effect of retrieval on downstream tasks. In this subsection, we  
 1158 continue the investigation into the downstream performance on the theorem-based tasks—TheoremQA  
 1159 Questions, TheoremQA Theorems, and AoPS—by evaluating gpt-4o-2024-08-06 on these  
 1160 tasks. We run with with three types of documents in context: none (closed-book), random (randomly  
 1161 sampled document from the corpus), and oracle (the annotated corpus). We use a temperature of 0.2  
 1162 and top-p of 0.9, and average the results across 5 different random seeds. We show the results in  
 1163 48, and find that the oracle document can consistently improve the performance across all datasets.  
 1164 This improvement is statistically significant in TheoremQA Questions and TheoremQA Theorems,  
 1165 suggesting the benefits of using strong retrieval results.

## 1166 E ANNOTATOR INSTRUCTIONS

1167  
 1168 In this section, we describe the instructions for annotators to collect data in **BRIGHT**.

### 1169 E.1 STACKEXCHANGE

- 1170 1. Browse posts from the newest to the oldest.
- 1171 2. Discard posts without an answer accepted by the user or obtains more than 5 votes
- 1172 3. Discard answers of posts without URL links.
- 1173 4. For each link in the answer, write down the answers to: (1). why are the document and the  
 1174 post relevant; (2). what is the reasoning required to understand the relevance between the  
 1175 post and the document. If these are not possible, discard the link.
- 1176 5. Check whether the linked documents provide critical information to understand the post or  
 1177 address the questions. The relevance could include explaining critical concepts or details or  
 1178 providing theorems, lemmas or code pieces that would contribute to solving the problems.  
 1179 Refer to examples for a better understanding on the relevance (Examples in Table 20 to 26)

1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187 <sup>17</sup>[https://huggingface.co/datasets/code\\_search\\_net](https://huggingface.co/datasets/code_search_net)

- 1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199
6. Use LLMs (e.g., ChatGPT, Claude, etc.) to generate post keywords, or use the post title to search for web pages with large keyword or semantic overlap in Google. Search for at most 5 negative web pages per query. Ensure that hard negatives are totally unhelpful in addressing the post. Common hard negatives include descriptions or documentation of similar contexts but different problems, or code pieces that share the same variable names but are unrelated to the post.
  7. Check all the other negative passages, which include positive and negative documents in previously annotated examples in the same dataset. Ensure that all those documents and current problems focus on different sub-topics, and thus totally irrelevant.
  8. Split every web page into small passages either by two newline symbols, "#" in markdown files or fixed-length tokens, and fine-grain positive ones following criteria in step 4 and 5.

## 1200 E.2 THEOREMQA

1201  
1202 In TheoremQA, the main task for the annotator is to check if the GPT-4 rewritten questions are valid.  
1203 The specific instructions are as follows:

- 1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212
1. Read the rewritten question and determine if it is solvable.
  2. If it is solvable, read the original question and solution, and determine if the rewritten question is consistent with the original question. That is, the same reasoning steps and the final answer should hold.
  3. If it is also consistent, mark the question as valid, and make any minor edits to the problem statement (e.g., to improve grammar or fluency) as you see fit.
  4. If it is not solvable or not consistent, read the original question and solution, and correct the rewritten question if possible. If not, then discard the problem.

## 1213 E.3 AoPS

1214  
1215 In AoPS, annotators are tasked to find questions from the AoPS Wiki and record the problems:

- 1216  
1217  
1218  
1219  
1220  
1221  
1222
1. Browse through the AoPS Wiki and find topic/category pages ([example 1](#), [example 2](#)).
  2. Look through each page and find pages specific theorems or techniques that can be used to solve problems. The page should link to at least two competition problems ([example 1](#), [example 2](#)).
  3. Record the links of both the theorem/technique as well as the problem pages.

1223 The annotators are assigned a category to look for theorems in to avoid overlaps, and the categories  
1224 are {algebra, geometry, calculus, probability, number theory, other}. After all  
1225 links are collected, we use a web scraper to collect the problem statement and solutions, and we  
1226 manually check the quality of the scraped data.

## 1227 E.4 LEETCODE

1228  
1229 In LeetCode, annotators determine whether a question is grounded in real-world concepts. We give a  
1230 similar instruction to the annotator as to GPT-4:

- 1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241
1. Read the problem statement carefully.
  2. Categorize the question into one of three categories:
    - 0: The question is not grounded in any real-world concepts. The description only uses coding-specific terms, such as "linked list", "binary search", "palindrome", "sorting", etc..
    - 1: The question is not grounded in any real-world concepts or real-world concepts that are commonly used in the context of coding, such as needle in a haystack, strings/words, or a spiral matrix.
    - 2: The question is grounded in real-world concepts that are not commonly used in the context of coding, such as building height, planting trees, or games. It may still use some code-specific terms to specify the data structure involved.



## 1242 E.5 LLM USAGE

1243  
1244 There is no specific procedure for the annotators to use LLMs. The LLMs serve a tool to help  
1245 annotators understand queries and documents, i.e., whenever they fail to understand something,  
1246 they ask LLMs for clarification, explanation, etc. They are also used to summarize the content of  
1247 StackExchange posts for searching negative documents in Google for StackExchange. To diversify  
1248 the search results, the annotator prompts LLMs with role-playing scenarios (e.g., as a biology student)  
1249 to generate keywords of the posts. The responses from LLMs are not included in any of **BRIGHT**  
1250 datasets. Since the responses from LLMs are not guaranteed to be correct, the annotators always  
1251 search for trustworthy sources to verify the information. The LLMs used in the annotation include  
1252 ChatGPT<sup>18</sup>, GPT-4, and Claude-3.

## 1253 E.6 SENSITIVE INFORMATION

1254  
1255 All the data in **BRIGHT** are manually collected, carefully verified, and reviewed to remove any  
1256 personally identifiable information or offensive content.  
1257

## 1258 F LIMITATIONS AND FUTURE WORK

1259  
1260 One limitation of this work is that the judgment about relevance between queries and documents is  
1261 subjective. Even if the StackExchange answers are accepted by the users or obtain high votes, it is  
1262 not guaranteed that everyone will agree the referenced documents are relevant. We may not expect  
1263 human retrieval results to be exactly the same as our annotation. However, we mitigate this issues  
1264 by using multiple annotators and only retain the queries in which all annotators, including domain  
1265 experts, agree on the relevance. Therefore, we believe that the relevance judgment in **BRIGHT** is  
1266 reliable and consistent.

1267 Aside from retrieval, other embedding tasks such as clustering may also require reasoning. We have  
1268 not addressed multi-modal settings in this paper, but they represent intriguing avenues for future  
1269 exploration.  
1270

## 1271 G POTENTIAL SOCIAL IMPACTS

1272  
1273 This paper presents a new retrieval benchmark that features relevance beyond lexical and semantic  
1274 similarity and requires intensive reasoning to solve. There are many potential societal consequences  
1275 of our work, e.g., improving search algorithms, fostering better information access, developing more  
1276 advanced retrieval models, etc. It could also promote collaboration among researchers and facilitate  
1277 the development of more effective search engines, ultimately benefiting society by enhancing the way  
1278 people find and access information.  
1279

## 1280 H COMPARISON TO RAR-B

1281  
1282 While both the *Reasoning as Retrieval Benchmark* (RAR-b) and *BRIGHT* evaluate retrieval systems  
1283 on their reasoning abilities, they differ significantly in their approach and objectives:  
1284

- 1285 • **Dataset Construction:**

- 1286 – *RAR-b*: Adapts existing multiple-choice benchmarks into a retrieval format, where the  
1287 queries are original questions and the corpus consists of unique options from all the  
1288 questions.
- 1289 – *BRIGHT*: Purposefully built as a retrieval benchmark which uses queries and documents  
1290 in realistic retrieval scenarios.

- 1291 • **Document Characteristics:**

- 1292 – *RAR-b*: Often uses very short sequences (<10 words) derived from multiple-choice  
1293 options.  
1294

1295  
<sup>18</sup><https://chatgpt.com/>

- 1296
- 1297
  - *BRIGHT*: Focuses on substantially longer documents (>100 tokens), more closely
  - 1298 mirroring practical retrieval scenarios.
- 1299
  - **Practical Relevance:**
- 1300
  - *RAR-b*: It’s more conceptual than practical, as real-world retrieval typically involves
  - 1301 retrieving documents instead of answers, making this test more about abstract reasoning
  - 1302 abilities.
- 1303
  - *BRIGHT*: Designed to reflect real-world information seeking behaviors and needs.

1304 We believe both datasets effectively assess retrievers’ capabilities in handling reasoning-intensive  
1305 queries, and we’ll be adding clearer comparisons of RAR-b in our revision.  
1306

1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

1350 Table 13: **Prompt used to rewrite TheoremQA questions.** Blue text denote instance-specific inputs.  
 1351 Violet text denotes the rewritten question and answer outputted by GPT-4. For each question, we  
 1352 provide the original question and answer, as well as the theorem name and theorem definition from  
 1353 the original TheoremQA dataset. The instruction prompts the model to rewrite the question with a  
 1354 different surface form without changing the reasoning steps or the final answer. We hand-write three  
 1355 examples to illustrate the rewriting process. We also allow the model to skip the question.

---

1356 Rewrite the following question such that the logical steps and final  
 1357 answer are still the same, only change the surface form of the question.  
 1358 Avoid using the words related to the theorem used in the question. You  
 1359 can do this by using real-world examples and applications to illustrate  
 1360 the concepts, in a way that is easier to understand for a layman. Try  
 1361 to ground the question in a real-world context such as finance and  
 1362 engineering problems, but be creative and feel free to use any domain  
 1363 you like! However, the question should be standalone and solvable.  
 1364 Rewrite the question in a json format, with the fields "question" and  
 1365 "answer". The "theorem" field indicates the theorem used in the  
 1366 question, but you should not use the words related to the theorem in the  
 1367 rewritten question. If you do not think that the question cannot be  
 1368 rewritten in layman's terms in a standalone fashion, then you can write  
 1369 a json object with the field "skip" set to True.

```

1369 { "question": "In a group of 10 people, each of whom has one of 3
1370 different eye colors, at least how many people must have the same eye
1371 color?", "answer": "4", "theorem": "pigeonhole principle",
1372 "theorem_definition": "The Pigeonhole Principle is a fundamental
1373 concept in combinatorics, a branch of mathematics that deals with
1374 counting and arranging objects. [additional text omitted...]" }
1375 one demonstration was omitted...
1376 { "question": "dy/dt = \sqrt{t}, y(1) = 1. What is y(4)?", "answer":
1377 "5.667", "theorem": "integral rules", "theorem_definition": "Integral
1378 rules in calculus are a set of techniques [additional text omitted...]" }
1379 { "question": "You're tracking the growth of a plant from a seed. The
1380 rate at which the plant grows in height is equal to the square root of
1381 the number of days since you planted it. One day after the seed was
1382 first planted, the plant was 1 inch tall. How tall will the plant be
1383 after 4 days?", "answer": "5.667" }
1384 { "question": "Given image \begin{tabular}{|l|l|l|} \hline 7 & 1 & 6 & 0 \\
1385 \ 3 & 3 & 7 & 6 \\ \ 6 & 6 & 5 & 7 \\ \hline \end{tabular}, and the bit
1386 -depth of the image is 4. Suppose you want to use the thresholding
1387 technique to segment the image. What is the appropriate threshold value
1388 based on the histogram of the image? Follow the following rule when you
1389 do thresholding or grouping: pixel $(i, j)\in$ Group A pixels if $g(i, j)
1390 \leq$ current threshold $\mathit{T}$; pixel $(i, j)\in$ Group B pixels
1391 otherwise, where $g(i, j)$ is the intensity value of pixel $(i, j)$.",
1392 "answer": "4", "theorem": "image contrast ", "theorem_definition":
1393 "Image contrast in signal processing refers to the difference in
1394 intensity or color between various elements or regions within an image.
1395 [additional text omitted...]" }
1396 { "question": "Imagine you have a collection of 12 unique stickers,
1397 each with a number from 0 to 7 representing its rarity level. You're
1398 planning to divide these stickers into two groups based on their rarity
1399 to make it easier for collectors to understand. You decide to use a
1400 method where stickers with a rarity level at or below a certain number
1401 go into the 'Common' group, and those above this number go into the
1402 'Rare' group. Given the distribution of stickers' rarity levels as
1403 follows: two stickers each of levels 7, 6, and 5; three stickers of
1404 level 3; and one sticker each of levels 1 and 0. What is the rarity
1405 level that should be used as the cutoff to divide the stickers into the
1406 'Common' and 'Rare' groups, ensuring a balanced understanding of
1407 rarity?", "answer": "4" }

```

---

1404 Table 14: **TheoremQA-Q Theorem names and the excluded topic names from CAMEL-Math.**

1405	TheoremQA-Q Theorem	CAMEL-Math Topics
1406	acyclic graph	Graph theory, Combinatorics
1407	binomial theorem	Combinatorics, Algebra
1408	catalan-mingantu number	Combinatorics
1409	cauchy’s integral theorem	Complex analysis
1410	cayley’s formula	Graph theory, Combinatorics
1411	convexity	Optimization
1412	cramer rao lower bound	Statistics
1413	definite matrix criteria	Algebra, Linear algebra
1414	derivative chain rule	Calculus
1415	double integral theorem	Calculus
1416	eigenvalues and eigenvectors	Linear algebra
1417	euler’s method	Calculus, Numerical analysis, Differential equations
1418	euler’s theory	Graph theory, Combinatorics
1419	expected utility	Game theory
1420	fisher information	Statistics
1421	fourier’s theorem	Fourier analysis
1422	gauss’s lemma	Number theory, Algebra
1423	integral rules	Calculus
1424	intermediate value theorem	Calculus
1425	isomorphisms	Group theory, Algebra
1426	jensen’s inequality	Statistics, Probability
1427	l’hopital’s rule	Calculus
1428	limit laws for sequences	Calculus
1429	limiting theorem	Calculus
1430	linear independence	Algebra, Linear algebra
1431	linear subspaces	Algebra, Linear algebra
1432	linear systems	Algebra, Linear algebra
1433	liouville’s theorem	Complex analysis
1434	martingale	Statistics, Probability
1435	matrix determinant formula	Algebra, Linear algebra
1436	maximal planar graph	Graph theory, Combinatorics
1437	maximum entropy	Statistics, Probability
1438	message passing algorithm	Graph theory, Combinatorics
1439	multinomial theorem	Combinatorics
1440	newton-raphson method	Calculus, Numerical analysis
1441	order	Group theory, Algebra
1442	ordinary differential equation	Calculus, Differential equations
1443	p-value	Statistics
1444	pigeonhole principle	Combinatorics
1445	poisson process	Statistics, Probability
1446	projection theory	Algebra, Linear algebra
1447	ramsey’s theorem	Graph theory, Combinatorics
1448	rolle’s theorem	Calculus
1449	series convergence	Calculus
1450	shortest path	Graph theory, Optimization, Combinatorics
1451	similarity	Geometry
1452	squeeze theorem	Calculus
1453	stirling number of the first kind	Combinatorics
1454	stirling number of the second kind	Combinatorics
1455	t-test	Statistics
1456	taylor’s approximation theorem	Calculus
1457	trapezoidal rule	Calculus, Numerical analysis
	vertex cover	Graph theory, Combinatorics
	viterbi algorithm	Statistics
	wave theorem	Differential equations

1458  
 1459  
 1460  
 1461  
 1462  
 1463  
 1464  
 1465  
 1466  
 1467  
 1468  
 1469  
 1470  
 1471  
 1472  
 1473  
 1474  
 1475  
 1476  
 1477  
 1478  
 1479  
 1480  
 1481  
 1482  
 1483  
 1484  
 1485  
 1486  
 1487  
 1488  
 1489  
 1490  
 1491  
 1492  
 1493  
 1494  
 1495  
 1496  
 1497  
 1498  
 1499  
 1500  
 1501  
 1502  
 1503  
 1504  
 1505  
 1506  
 1507  
 1508  
 1509  
 1510  
 1511

Table 15: Subfield name in the original TheoremQA dataset and their corresponding categories in BRIGHT. Each category is used to specify the excluded problem-solution pairs from MATH and AoPS.

Name	Category
Algebra	algebra
Atomic physics	others
Calculus	calculus
Celestial mechanics	others
Classic mechanics	others
Combinatorics	number theory
Complex analysis	calculus
Computer networking	others
Condensed matter physics	others
Derivatives	calculus
Economics	others
Electromagnetism	others
Equity investments	others
Fixed income	others
Fluid mechanics	others
Functional analysis	calculus
Geometry	geometry
Graph theory	others
Group theory	algebra
Information theory	others
Kinetics	others
Machine learning	others
Mathematical analysis	calculus
Number theory	number theory
Numerical analysis	calculus
Optics	others
Particle	others
Portfolio management	others
Probability theory	probability
Quantitive methods	others
Quantum	others
Real analysis	calculus
Relativity	others
Signal processing	others
Statistical physics	others
Statistics	others
Stochastic process	probability
Thermodynamics	others
Wave	others



1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531  
 1532  
 1533  
 1534  
 1535  
 1536  
 1537  
 1538  
 1539  
 1540  
 1541  
 1542  
 1543  
 1544  
 1545  
 1546  
 1547  
 1548  
 1549  
 1550  
 1551  
 1552  
 1553  
 1554  
 1555  
 1556  
 1557  
 1558  
 1559  
 1560  
 1561  
 1562  
 1563  
 1564  
 1565

Table 16: **AoPS theorems and techniques and the excluded topic names from CAMEL-Math and TheoremQA.**

AoPS Theorem	CAMEL-Math Topics	TheoremQA Theorems
Binomial Theorem	Combinatorics, Algebra	binomial theorem, multinomial theorem
Vietas Formulas	Algebra	vieta’s formula, birg-vieta’s theorem
Ptolemy’s theorem	Geometry	properties of kites, similarity, rhombus, rectangle, quadrilateral, triangle, isosceles triangle, parallelogram
Recursive Series	Calculus, Algebra	
Power of a Point	Geometry	similarity
Ball and urn	Combinatorics	
Newtons Sums	Algebra	vieta’s formula
Probability	Probability, Statistics	probability
Fibonacci sequence	Combinatorics, Number theory	
Chicken McNugget Theorem	Combinatorics, Number theory	
Central Tendency	Statistics	
Principle of Inclusion Exclusion	Combinatorics, Set theory	inclusion-exclusion principle
Factorial	Combinatorics	
Picks Theorem	Geometry	similarity, triangle, isosceles triangle, parallelogram, rhombus, quadrilateral, rectangle, triangle midsegment theorem
Shoelace Theorem	Geometry	rhombus, rectangle, quadrilateral, triangle, isosceles triangle, parallelogram
Geometric probability	Geometry, Probability	
Euclidean algorithm	Number theory, Cryptography, Algebra	euclidean algorithm
Mass Points	Geometry	similarity, triangle, isosceles triangle, parallelogram, rhombus, quadrilateral, rectangle, triangle midsegment theorem
Geometric Series	Algebra	
Triangle Inequality	Geometry	inequalities, triangle, isosceles triangle
Simons Favorite Factoring Trick	Number theory, Algebra	
Properties of Logarithms	Algebra	
Fermat’s Little Theorem	Number theory, Cryptography	fermat’s little theorem, euler’s totient theorem

1566

1567

Table 17: **Prompt used to check TheoremQA-Theorem documents.** Blue text denote instance-specific inputs. For each question, we provide the problem statement, the solution, the theorem name, the theorem definition, and the document text from a ProofWiki document.

1569

1570

1571

1572

1573

1574

1575

1576

1577

---

Instruction: Determine if the given text can be helpful in understanding and solving the given problem. Use the theorem, its definition, and the problem solution to help you make a decision. The text can be helpful if it uses very similar reasoning steps as the solution and applies the theorem in a related way as the solution applies the theorem to solve the problem. The text should be able to assist a student who is not familiar with the theorem in ultimately solving the problem.

1578

1579

1580

1581

1582

1583

1584

1585

1586

The input is given to you in a json format with the following keys:  
 Problem statement: The problem statement that the student is trying to solve.

Solution: The solution to the problem.

Theorem: The theorem that is used in the solution.

Theorem definition: The definition of the theorem.

Text: The text that you need to evaluate.

Think step by step and reason about the theorem and the text first before finally making a decision. Output True if the text is helpful, and False if it is not.

1587

1588

1589

1590

1591

```
{ "Problem statement": "", "Solution": "", "Theorem": "", "Theorem
definition": "", "Text": "", }
```

Now, write your answer in the following format:

Reasoning: [your reasoning here]

Answer: [True/False]

---

1592

1593

1594

1595

Table 18: **Theorems and techniques used in the AoPS dataset, and their corresponding subfield categories.**

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

Name	Category
Ball and urn	number theory
Binomial Theorem	algebra
Central Tendency	others
Chicken McNugget Theorem	algebra
Euclidean algorithm	number theory
Factorial	number theory
Fermat's Little Theorem	number theory
Fibonacci sequence	number theory
Geometric Series	algebra
Geometric probability	probability
Mass Points	geometry
Newtons Sums	algebra
Picks Theorem	geometry
Power of a Point	geometry
Principle of Inclusion Exclusion	number theory
Probability	probability
Properties of Logarithms	algebra
Ptolemy's theorem	geometry
Recursive Series	algebra
Shoelace Theorem	geometry
Simon's Favorite Factoring Trick	number theory
Triangle Inequality	algebra
Vieta's Formulas	algebra

1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673

**Table 19: Prompt used to check LeetCode questions.** Blue text denote instance-specific inputs. Violet text denotes the output from GPT-4. For each question, we provide the title and the problem statement. We use six in-context demonstrations, and the instruction prompts the model to categorize the question based on the criteria provided. Although there are three possible labels, we only keep questions that were labeled as 2, which are grounded in real-world concepts. For sake of brevity, we only show one example, but other examples can be found on the code repo.

---

```

Read the coding question and categorize it using the following criteria:
0: The question is not grounded in any real-world concepts. The
description only uses coding-specific terms, such as "linked list",
"binary search", "palindrome", "sorting", etc..
1: The question is not grounded in any real-world concepts or
real-world concepts that are commonly used in the context of coding,
such as needle in a haystack, strings/words, or a spiral matrix.
2: The question is grounded in real-world concepts that are not
commonly used in the context of coding, such as building height,
planting trees, or games. It may still uses some code-specific terms to
specify the data structure involved.

You should only consider the initial problem statement and problem
title, not the examples or constraints.
Use the following examples to help you categorize the question:

Example 1:
{{
"title": "Merge Two Sorted Lists",
"question": "You are given the heads of two sorted linked lists `list1`
and `list2`.
Merge the two lists in a one sorted list. The list should be made
by splicing together the nodes of the first two lists.

Return the head of the merged linked list."

Example 1:
rest of the example omitted...
}}

{{
"label": 0
}}

Examples 2-6 are omitted for brevity...

Now, consider the question below and categorize it using the criteria
above. Output your answer in a json format:
{{
"title": "Container With Most Water",
"question": "You are given an integer array `height` of length `n`.
There are `n` vertical lines drawn such that the two endpoints of the
`ith` line are `(i, 0)` and `(i, height[i])`.

Find two lines that together with the x-axis form a container, such that
the container contains the most water.

Return the maximum amount of water a container can store.
rest of the question omitted...
}}

{ "label": 2 }

```

---

1674 Table 20: **Biology example.** The positive document explains the concept of meristem, which is the  
 1675 mechanism by which plant cells develop into other tissues and organs. This provides critical support  
 1676 for why a tree trunk sprouts and grows after being cut.

1677	
1678	<i>Query</i>
1679	After I cut trees into logs and remove the branches in winter, they start growing. They sprout
1680	out and grow completely normal looking stems and leaves and maintain them all summer. The
1681	sprouts mostly appear around the cut branches. Sometimes they last all winter and grow for
1682	another year. How does a tree trunk sprout and grow after being cut?
1683	
1684	<i>Chain-of-thought reasoning to find documents</i>
1685	The essential question is to figure out the growing mechanism of trees after being cut. It is
1686	usually related to cell division, which later develops into tissues and organs. We need to find
1687	relevant information of these cells in plants and their growing capabilities after being cut.
1688	<i>Example positive document</i>
1689	Meristem
1690	Tunica-carpus model of the apical meristem (growing tip). The epidermal (L1) and subepider-
1691	mal (L2) layers form the outer layers called the tunica. The inner L3 layer is called the corpus.
1692	Cells in the L1 and L2 layers divide in a sideways fashion, which keeps these layers distinct,
1693	whereas the L3 layer divides in a more random fashion.
1694	In cell biology, the meristem is a type of tissue found in plants. It consists of undifferentiated
1695	cells (meristematic cells) capable of cell division. Cells in the meristem can develop into all
1696	the other tissues and organs that occur in plants. These cells continue to divide until they
1697	become differentiated and lose the ability to divide.
1698	There are three types of meristematic tissues: apical (at the tips), intercalary or basal (in
1699	the middle), and lateral (at the sides also known as cambium). At the meristem summit,
1700	there is a small group of slowly dividing cells, which is commonly called the central zone.
1701	Cells of this zone have a stem cell function and are essential for meristem maintenance. The
1702	proliferation and growth rates at the meristem summit usually differ considerably from those
1703	at the periphery.
1704	...
1705	Under appropriate conditions, each shoot meristem can develop into a complete, new plant or
1706	clone. Such new plants can be grown from shoot cuttings that contain an apical meristem.
1707	...
1708	<i>Example negative document</i>
1709	Cutting firewood from felled trees safely and easily (bucking)
1710	There are loads of articles on the web showing how (and sometimes how not to!) cut down
1711	trees, but not so much about actually cutting firewood after the tree has been felled and stripped
1712	of its branches (called snedding or limbing in arborist speak). This ‘how to’ shows the easiest
1713	way to produce firewood by cutting many logs at the same time. I actually use my firewood
1714	cutting operation as part of my fitness regime, much more useful than going to the gym and
1715	the same muscle burn the day after!
1716	I start by cutting the trees into 10’ (3m) logs so that I can easily drag them to my cleared
1717	working area close to the pick up point.
1718	You can make a homemade ‘cradle’ that holds a whole bunch of logs from sturdy construction
1719	lumber (2 by 4). This enables you to safely cut many at once, as long as you put some of
1720	the heavy logs both at the bottom and on the top. Don’t forget to make the cradle narrow
1721	enough for the size of chainsaw bar you have. Usually not more than 12” or 300mm inside
1722	measurement. Make sure that it is constructed in such a way that you cannot hit any nails etc
1723	with the chainsaw.
1724	Place the logs onto the cradle keeping the ends off the logs flush one side ( or at least staggered
1725	in 12” steps). If you are fussy like me you can mark one of the top logs at 12” (300mm)
1726	intervals with chalk to ensure even log lengths. Then cut alternate sides (so that it doesn’t
1727	topple over to one side), finishing off with a couple of cuts over the cradle itself, don’t forget
	to stop before you saw the thing in half lol! (See upgraded version below...)
	...

Table 21: **Earth Science example.** The positive document describes the Airy’s isostasy model, which is needed in calculating the hydrostatic equilibrium.

1728  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739  
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781

---

**Query**

---

How to calculate hydrostatic equilibrium?

I’m trying to solve the following problem. The sea level in the past was 200 m higher than today. The seawater became in isostatic equilibrium with the ocean basin. what is the increase in the depth  $x$  of the ocean basins? Water density is  $p_w = 1000 \text{ kg } m^{-3}$ , and mantle density is  $3300 \text{ kg } m^{-3}$

Using the compensation column, I reach:

$$x = (p_w * 200m)/3300 = 60.60 \text{ m}$$

but normally I expected to find 290 m.

Can someone explain to me what’s wrong?

---

**Chain-of-thought reasoning to find documents**

---

This problem is about the hydrostatic equilibrium. This may be related to Pascal’s law, which gives insights on the pressure change in fluid’s mechanics. To understand hydrostatic equilibrium or Pascal’s law, it would be helpful to provide a detailed explanation of the concept and formula used in calculations. Other content on the application of the theorems, or related lemmas that also focus on the discussion of isotatic equilibrium with the ocean basin would also be useful.

---

**Example positive document**

---

Airy

The basis of the model is Pascal’s law, and particularly its consequence that, within a fluid in static equilibrium, the hydrostatic pressure is the same on every point at the same elevation (surface of hydrostatic compensation):

$$h_1 \cdot p_1 = h_2 \cdot p_2 = h_3 \cdot p_3 = \dots h_n \cdot p_n$$

For the simplified picture shown, the depth of the mountain belt roots ( $b_1$ ) is calculated as follows:

$$(h_1 + c + b_1)p_c = (cp_c) + (b_1p_m)$$

$$b_1(p_m - p_c) = h_1p_c$$

$$b_1 = h_1p_c/(p_m - p_c)$$

where  $p_m$  is the density of the mantle (ca.  $3,300 \text{ kg } m^{-3}$ ) and  $p_c$  is the density of the crust (ca.  $2,750 \text{ kg } m^{-3}$ ). Thus, generally:

$$b_1 = 5h_1$$

...

---

**Example negative document**

---

Global or eustatic sea level has fluctuated significantly over Earth’s history. The main factors affecting sea level are the amount and volume of available water and the shape and volume of the ocean basins. The primary influences on water volume are the temperature of the seawater, which affects density, and the amounts of water retained in other reservoirs like rivers, aquifers, lakes, glaciers, polar ice caps and sea ice. Over geological timescales, changes in the shape of the oceanic basins and in land/sea distribution affect sea level. In addition to eustatic changes, local changes in sea level are caused by the earth’s crust uplift and subsidence.

Over geologic time sea level has fluctuated by more than 300 metres, possibly more than 400 metres. The main reasons for sea level fluctuations in the last 15 million years are the Antarctic ice sheet and Antarctic post-glacial rebound during warm periods.

The current sea level is about 130 metres higher than the historical minimum. Historically low levels were reached during the Last Glacial Maximum (LGM), about 20,000 years ago. The last time the sea level was higher than today was during the Eemian, about 130,000 years ago.

...

---

1782 Table 22: **Economics example**. The positive document describes the stochastic dominance, which is  
 1783 the underlying concept used to solve the problem.

1784	
1785	<b>Query</b>
1786	Which of these two lotteries, a consumer with Von-Neumann Morgenstern preferences will
1787	choose under exponential distribution?
1788	Consider two lotteries each having an exponential distribution. The function of cumulative
1789	distribution of an exponential distribution is:
1790	$F(x; \lambda) = 1 - e^{-\lambda x} \forall x \in R_+$ .
1791	The expected gain given by this distribution is $E[X] = 1/\lambda$ .
1792	Suppose that the first lottery, $F_1$ has a parameter $\lambda_1$ and the second, $F_2$ has a parameter
1793	$\lambda_2$ . Suppose that $\lambda_1 < \lambda_2$ . Which of these two lotteries, a consumer with Von-Neumann
1794	Morgenstern preferences will choose?
1795	I am confused on this question.
1796	<b>Chain-of-thought reasoning to find documents</b>
1797	The expected value is different from the expected utility. The preferences under uncertainty in
1798	this case should be linked to first-order stochastic dominance. We need to look up the detailed
1799	definition to understand how the problem could be solved by using that framework.
1800	
1801	<b>Example positive document</b>
1802	Statewise dominance implies first-order stochastic dominance (FSD), which is defined as:
1803	Random variable A has first-order stochastic dominance over random variable B if for any
1804	outcome x, A gives at least as high a probability of receiving at least x as does B, and for some
1805	x, A gives a higher probability of receiving at least x. In notation form, $P[A \geq x] \geq P[B \geq x]$
1806	for all x, and for some x, $P[A \geq x] > P[B \geq x]$ .
1807	In terms of the cumulative distribution functions of the two random variables, A dominating B
1808	means that $F_A(x) \leq F_B(x)$ for all x, with strict inequality at some x.
1809	In the case of non-intersecting distribution functions, the Wilcoxon rank-sum test tests for
1810	first-order stochastic dominance.
1811	...
1812	Extended example
1813	Consider three gambles over a single toss of a fair six-sided die:
1814	State (die result) 1 2 3 4 5 6
1815	gamble A wins \$ 1 1 2 2 2 2
1816	gamble B wins \$ 1 1 1 2 2 2
1817	gamble C wins \$ 3 3 3 1 1 1
1818	Gamble A statewise dominates gamble B because A gives at least as good a yield in all
1819	possible states (outcomes of the die roll) and gives a strictly better yield in one of them (state
1820	3). Since A statewise dominates B, it also first-order dominates B.
1821	...
1822	<b>Example negative document</b>
1823	Which Lottery Is Easiest To Win Within the U.S.?
1824	Winning the lottery is one of the least likely things to happen in your life. However, some
1825	lotteries are much easier to win than others.
1826	It's common for a lot of Americans to daydream about winning the lottery. These dreams
1827	became a reality for a few lucky winners as they became millionaires overnight. The cold
1828	truth is that the overwhelming majority of lottery dreams will remain dreams.
1829	Most people know that the odds of winning the lottery are not in your favor. If it were easy to
1830	win, everyone would play, and the prizes would only be a few cents.
1831	You're likelier to be struck by lightning, live to be 110 years old, and be declared a saint by
1832	the Catholic Church.
1833	The combination of high jackpots and low odds has caused many people to try and cheat in the
1834	lottery. The attempts might be successful for a short period, but they always end with prison
1835	sentences.
	Fortunately, there are a few ways that you can legally improve your odds of winning. It will
	start by choosing games that have the best odds.
	...



Table 23: **Psychology example.** The positive document describes the critical equation needed to explain the confusion in the post.

1836	
1837	
1838	
1839	<b>Query</b>
1840	Why do fNIRS devices commonly use two different frequencies?
1841	One of the most common techniques used for functional neuroimaging nowadays is functional
1842	near infra-red spectroscopy (fun fact: IIRC Natalie Portman worked on a research paper
1843	involving fNIRS as the modality), which shines near infrared light into the brain from a source
1844	to a detector (both called optodes) in a "banana" shape.
1845	It's not uncommon to read that most of these devices, be they continuous wave (CW) or one
1846	of the two kinds that involve fast modulation, frequency domain (FD) or time domain (TD),
1847	require two separate frequencies to be emitted. For instance, NIRx explains it as follows on
1848	their website:
1849	"For neuro-imaging applications it is by far most common to illuminate with two discrete
1850	wavelength, which is the minimum requirement to assess relative variations of both oxygena-
1851	tion states of the hemoglobin molecule independently."
1852	Why is that the case?
1853	I haven't delved into the intricacies of it, but no reason immediately jumps out at me. For
1854	instance, in the case of CW, the relative difference in intensity is all that matters, so why do
1855	we need two frequencies?
1856	<b>Chain-of-thought reasoning</b>
1857	In order to understand why we need two frequencies in fNIRS, we will need to check the
1858	measurements and calculation in continuous wave systems. This may be related to oxy-
1859	genated (HbO) and deoxygenated (HbR) hemoglobin in the brain. We need to refer to the
1860	dual-wavelength approach, which is fundamental to fNIRS's ability to provide meaningful
1861	information about local brain activity.
1862	<b>Example positive document</b>
1863	...
1864	Where OD is the optical density or attenuation, $I_0$ is emitted light intensity, $I$ is measured
1865	light intensity, $\epsilon$ is the attenuation coefficient, $[X]$ is the chromophore concentration, $l$ is the
1866	distance between source and detector and $DPF$ is the differential path length factor, and $G$ is
1867	a geometric factor associated with scattering.
1868	When the attenuation coefficients $\epsilon$ are known, constant scattering loss is assumed, and the
1869	measurements are treated differentially in time, the equation reduces to:
1870	$\Delta[X] = \Delta OD / (\epsilon d)$
1871	Where $d$ is the total corrected photon path-length.
1872	Using a dual wavelength system, measurements for HbO <sub>2</sub> and Hb can be solved from the
1873	matrix equation:
1874	$\begin{pmatrix} \Delta OD_{\lambda_1} \\ \Delta OD_{\lambda_2} \end{pmatrix} = \begin{pmatrix} \epsilon_{\lambda_1}^{Hb} d & \epsilon_{\lambda_1}^{HbO_2} d \\ \epsilon_{\lambda_2}^{Hb} d & \epsilon_{\lambda_2}^{HbO_2} d \end{pmatrix} \begin{pmatrix} \Delta[X]^{Hb} \\ \Delta[X]^{HbO_2} \end{pmatrix}$
1875	...
1876	<b>Example negative document</b>
1877	What Is Wave Frequency?
1878	The number of waves that pass a fixed point in a given amount of time is wave frequency.
1879	Wave frequency can be measured by counting the number of crests (high points) of waves
1880	that pass the fixed point in 1 second or some other time period. The higher the number is,
1881	the greater the frequency of the waves. The SI unit for wave frequency is the hertz (Hz), where 1
1882	hertz equals 1 wave passing a fixed point in 1 second. The Figure below shows high-frequency
1883	and low-frequency transverse waves.
1884	Q: The wavelength of a wave is the distance between corresponding points on adjacent waves.
1885	For example, it is the distance between two adjacent crests in the transverse waves in the
1886	diagram. Infer how wave frequency is related to wavelength.
1887	A: Waves with a higher frequency have crests that are closer together, so higher frequency
1888	waves have shorter wavelengths.
1889	...

Table 24: **Robotics example**. The positive document provides the code pieces to modify.

1890	
1891	
1892	
1893	<b>Query</b>
1894	Can't see debug messages using RCLCPP_DEBUG
1895	I can't see messages using RCLCPP_DEBUG by terminal and rqt, but I can using other levels
1896	of verbosity( INFO, ERROR, FATAL. . .).Selecting debug in rqt to see those messages doesn't
1897	work either.
1898	I'm using rolling, working in C++ in a plugin of a controller and launching
1899	tb3_simulation_launch.py from nav2_bringup.
1900	I also saw a post here where they recommended to set the environment variable:
1901	RCLCPP_LOG_MIN_SEVERITY=RCLCPP_LOG_MIN_SEVERITY_DEBUG
1902	but that didn't work either. It must be something silly that I'm missing. Has this ever happened
1903	to you?
1904	Thank you
1905	<b>Chain-of-thought reasoning</b>
1906	The problem looks like to be related to the setting of log level. If selecting debug in rqt does
1907	not work, we can check the code about potential arguments that could affect the logging
1908	messages and verbosity,
1909	<b>Example positive document</b>
1910	...
1911	declare_log_level_cmd = DeclareLaunchArgument(
1912	'log_level', default_value='info', description='log level'
1913	)
1914	load_nodes = GroupAction(
1915	condition=IfCondition(PythonExpression([not , use_composition])),
1916	actions=[
1917	SetParameter('use_sim_time', use_sim_time),
1918	Node(
1919	package='nav2_controller',
1920	executable='controller_server',
1921	output='screen',
1922	respawn=use_respawn,
1923	respawn_delay=2.0,
1924	parameters=[configured_params],
1925	arguments=['-ros-args', '-log-level', log_level],
1926	remappings=remappings + [('cmd_vel', 'cmd_vel_nav')],
1927	),
1928	...
1929	<b>Example negative document</b>
1930	The logger methods are named after the level or severity of the events they are used to track.
1931	The standard levels and their applicability are described below (in increasing order of severity):
1932	
1933	DEBUG
1934	Detailed information, typically of interest only when diagnosing problems.
1935	INFO
1936	Confirmation that things are working as expected.
1937	WARNING
1938	An indication that something unexpected happened, or indicative of some problem in the near
1939	future (e.g. 'disk space low'). The software is still working as expected.
1940	ERROR
1941	Due to a more serious problem, the software has not been able to perform some function.
1942	CRITICAL
1943	A serious error, indicating that the program itself may be unable to continue running.
	...

1944 Table 25: **Stack Overflow example**. The positive document provides the documentation that explains  
 1945 the method to solve the problem.

1946	
1947	<b>Query</b>
1948	Is there a Snowflake command that can do the following:
1949	
1950	a,b,c
1951	1,10,0.1
1952	2,11,0.12
1953	3,12,0.13
1954	to a table like this:
1955	
1956	key,value
1957	a,1
1958	a,2
1959	a,3
1960	b,10
1961	b,11
1962	b,13
1963	c,0.1
1964	c,0.12
1965	c,0.13
1966	?
1967	This operation is often called melt in other tabular systems, but the basic idea is to convert the
1968	table into a list of key value pairs.
1969	
1970	There is an UNPIVOT in SnowSQL, but as I understand it UNPIVOT requires to manually
1971	specify every single column. This doesn't seem practical for a large number of columns.
1972	<b>Chain-of-thought reasoning steps to find documents</b>
1973	This operation is a kind of table transformation like reshaping, serialization, or flattening. We
1974	are looking for an operation in Snowflake that can take 2-dimensional values, and transform
1975	them into an view that presents one-one correlation mapping between key and value.
1976	<b>Example positive document</b>
1977	FLATTEN
1978	Flattens (explodes) compound values into multiple rows.
1979	FLATTEN is a table function that takes a VARIANT, OBJECT, or ARRAY column and
1980	produces a lateral view (i.e. an inline view that contains correlation referring to other tables
1981	that precede it in the FROM clause).
1982	FLATTEN can be used to convert semi-structured data to a relational representation.
1983	Syntax
1984	FLATTEN( INPUT => <expr> [ , PATH => <constant_expr> ]
1985	[ , OUTER => TRUE   FALSE ]
1986	[ , RECURSIVE => TRUE   FALSE ]
1987	[ , MODE => 'OBJECT'   'ARRAY'   'BOTH' ] )
1988	...
1989	<b>Example negative document</b>
1990	SYSTEM\$EXTERNAL_TABLE_PIPE_STATUS
1991	Retrieves a JSON representation of the current refresh status for the internal (hidden) pipe
1992	object associated with an external table.
1993	Automatically refreshing the metadata for an external table relies internally on Snowpipe,
1994	which receives event notifications when changes occur in the monitored cloud storage. For
1995	more information, see Introduction to external tables.
1996	Syntax
1997	SYSTEM\$EXTERNAL_TABLE_PIPE_STATUS( '<external_table_name>' )
	...

1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009  
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051

Table 26: **Sustainable Living example.** The positive document explains the critical concept of soluble salts, which indicate that they contain dissolved minerals and could be harmful to plants.

<i>Query</i>	
	How good is it to reuse water from plant pots?
	I'm living in an apartment, and after I water my plants the water goes to plates below the pots. The pots are in a metallic structure above the plates, so I can take the plates to reuse the water (throwing it at the plants again).
	This reuse seems beneficial, because I think I can get rid of mosquitoes that would reproduce in the stagnated water. And also some nutrients of the soil (as well as earthworms) can return to the vase.
	Is there some negative points in doing that?
<i>Chain-of-thought reasoning to find documents</i>	
	The water that accumulates in the plates as a result of watering is likely to contain minerals, soluble salts, and other materials that exist in soil or plants. To figure out whether there is any negative point in reusing that water, we thus need to understand whether the components in that water will result in any adverse effects.
<i>Example positive document</i>	
	<b>Soluble Salts</b>
	Soluble salts may accumulate on the top of the soil, forming a yellow or white crust. A ring of salt deposits may form around the pot at the soil line or around the drainage hole. Salts may also build up on the outside of clay pots. In house plants, signs of excess soluble salts include reduced growth, brown leaf tips, dropping of lower leaves, small new growth, dead root tips, and wilting.
	Soluble salts are minerals dissolved in water. Fertilizer dissolved in water becomes a soluble salt. When water evaporates from the soil, the minerals or salts stay behind. As the salts in the soil become more and more concentrated, it becomes more difficult for plants to take up water. If salts build up to an extremely high level, water can be taken out of the root tips, causing them to die. High levels of soluble salts damage the roots directly, weakening the plant and making it more susceptible to attack from insects and diseases. One of the most common problems associated with high salt levels is root rot.
	The best way to prevent soluble salt injury is to stop the salts from building up. When watering, allow some water to drain through the container and then empty the saucer. Do not allow the pot to sit in water. If the drained water is absorbed by the soil, the salts that were washed out are reabsorbed through the drainage hole or directly through a clay pot.
	...
<i>Example negative document</i>	
	<b>Water reuse in California</b>
	Water reuse in California is the use of reclaimed water for beneficial use. As a heavily populated state in the drought-prone arid west, water reuse is developing as an integral part of water in California enabling both the economy and population to grow.
	Wastewater Reclaimed water is treated wastewater that comes from homes and businesses, such as sink water, shower water, and toilet water including everything dumped into wastewater drains from laundry soap to bleach to oil to human waste. Wastewater can be divided into greywater and blackwater, with the first being defined as water that had been used for laundry, bathing, sink washing, and dishwashers. Blackwater is defined as sewage that includes feces from toilets.[1] Due to the low amounts of physical pollutants in greywater, most of its contaminants are dissolved organic matter, which can be physically filtered and cleaned through various membranes, as well as through biological treatment methods.
	...

Table 27: **LeetCode example.** Both the query and the positive document uses two pointers in the solution. The negative example is retrieved by BM25.

2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105

---

### *Query*

---

Given ‘n’ non-negative integers representing an elevation map where the width of each bar is ‘1’, compute how much water it can trap after raining.

**Example 1:**

**Input:** height = [0, 1, 0, 2, 1, 0, 1, 3, 2, 1, 2, 1]

**Output:** 6

**Explanation:** The above elevation map (black section) is represented by array [0, 1, 0, 2, 1, 0, 1, 3, 2, 1, 2, 1]. In this case, 6 units of rain water (blue section) are being trapped.

**Example 2:**

**Input:** height = [4, 2, 0, 3, 2, 5]

**Output:** 9

**Constraints:**

\* ‘n == height.length’

\* ‘1 <= n <= 2 \* 10<sup>4</sup>’

\* ‘0 <= height[i] <= 10<sup>5</sup>’

---

### *Chain-of-thought reasoning to find documents*

---

This problem can be solved using a two-pointer approach, and uses ideas from dynamic programming to keep track of the maximum height to the left and right of each bar. We can find other example code that also use these techniques.

---

### *Example positive document*

---

def max\_area(height):

"""You are given an integer array ‘height’ of length ‘n’. There are ‘n’ vertical lines drawn such that the two endpoints of the ‘ith’ line are ‘(i, 0)’ and ‘(i, height[i])’.

Find two lines that together with the x-axis form a container, such that the container contains the most water.

Return the maximum amount of water a container can store.

**Notice** that you may not slant the container.

**Example 1:**

**Input:** height = [1, 8, 6, 2, 5, 4, 8, 3, 7]

**Output:** 49

**Constraints:**

\* ‘n == height.length’

\* ‘2 <= n <= 10<sup>5</sup>’

\* ‘0 <= height[i] <= 10<sup>4</sup>’"""

max\_area, left, right = 0, 0, len(height) - 1

while left < right:

max\_area = max(max\_area, min(height[left], height[right]) \* (right - left))

if height[left] < height[right]:

left += 1

else:

right -= 1

return max\_area

---

### *Example negative document*

---

def get\_power(x):

"""Your country has an infinite number of lakes. Initially, all the lakes are empty, but when it rains over the ‘nth’ lake, the ‘nth’ lake becomes full of water. If it rains over a lake that is **full of water**, there will be a **flood**. Your goal is to avoid floods in any lake.

Given an integer array ‘rains’ where: ...rest of the document omitted

---

Table 28: **Pony example**

2106	
2107	
2108	<b><i>Query</i></b>
2109	
2110	Given the lengths of a triangle’s sides, write a pony program to classify it as equilateral, isosceles or scalene.
2111	
2112	<b><i>Chain-of-thought reasoning to find documents</i></b>
2113	
2114	To determine the category of the triangle, we need to compare the side lengths. A common syntax for conditions and comparison used in many other programming languages is the ‘if’ or ‘switch’ related to control structure.
2115	
2116	We need to find the syntax that allows to specify such conditions.
2117	<b><i>Example positive document</i></b>
2118	
2119	# Control Structures
2120	
2121	To do real work in a program you have to be able to make decisions, iterate through collections of items and perform actions repeatedly. For this, you need control structures. Pony has control structures that will be familiar to programmers who have used most languages,
2122	
2123	used most languages,
2124	such as ‘if’, ‘while’ and ‘for’, but in Pony, they work slightly differently.
2125	
2126	## Conditionals
2127	
2128	The simplest control structure is the good old ‘if’. It allows you to perform some action only when a condition is true.
2129	In Pony it looks like this:
2130	
2131	if condition then
2132	control_body
2133	end
2134	
2135	Here is a simple example:
2136	
2137	if a > b then
2138	env.out.print("a is bigger")
2139	end
2140	
2141	Often the condition may be composed of many sub conditions connected by ‘and’ and ‘or’.
2142	...
2143	<b><i>Example negative document</i></b>
2144	
2145	# Classes
2146	
2147	Just like other object-oriented languages, Pony has __classes__.
2148	A class is declared with the keyword ‘class’, and it has to have a name that starts with a capital letter, like this:
2149	
2150	class Wombat
2151	
2152	Do all types start with a capital letter?
2153	Yes! And nothing else starts with a capital letter.
2154	So when you see a name in Pony code, you will instantly know whether it’s a type or not.
2155	## What goes in a class?
2156	
2157	A class is composed of:
2158	
2159	1. Fields.
	2. Constructors.
	3. Functions.
	...



2160  
2161  
2162  
2163  
2164  
2165  
2166  
2167  
2168  
2169  
2170  
2171  
2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187  
2188  
2189  
2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202  
2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213

Table 29: **AoPS example.** The negative example is retrieved by BM25, and it uses stars and bars to solve the problem, whereas the query and gold document use the Chicken McNugget Theorem.

---

*Query (from 2015 AMC 10B Problem 15)*

---

The town of Hamlet has 3 people for each horse, 4 sheep for each cow, and 3 ducks for each person. Which of the following could not possibly be the total number of people, horses, sheep, cows, and ducks in Hamlet?

(A) 41    (B) 47    (C) 59    (D) 61    (E) 66

---

*Chain-of-thought reasoning to find documents*

---

We can use the [Chicken McNugget Theorem](#) to solve this problem. We can find other solutions that also apply this theorem.

---

*Example positive document*

---

Find the sum of all positive integers  $n$  such that, given an unlimited supply of stamps of denominations 5,  $n$ , and  $n + 1$  cents, 91 cents is the greatest postage that cannot be formed. By the Chicken McNugget theorem, the least possible value of  $n$  such that 91 cents cannot be formed satisfies  $5n - (5 + n) = 91 \implies n = 24$ , so  $n$  must be at least 24.

For a value of  $n$  to work, we must not only be unable to form the value 91, but we must also be able to form the values 92 through 96, as with these five values, we can form any value greater than 96 by using additional 5 cent stamps.

Notice that we must form the value 96 without forming the value 91. If we use any 5 cent stamps when forming 96, we could simply remove one to get 91. This means that we must obtain the value 96 using only stamps of denominations  $n$  and  $n + 1$ .

Recalling that  $n \geq 24$ , we can easily figure out the working  $(n, n + 1)$  pairs that can be used to obtain 96, as we can use at most  $\frac{96}{24} = 4$  stamps without going over. The potential sets are  $(24, 25)$ ,  $(31, 32)$ ,  $(32, 33)$ ,  $(47, 48)$ ,  $(48, 49)$ ,  $(95, 96)$ , and  $(96, 97)$ .

The last two obviously do not work, since they are too large to form the values 92 through 94, and by a little testing, only  $(24, 25)$  and  $(47, 48)$  can form the necessary values, so  $n \in \{24, 47\}$ .  $24 + 47 = \boxed{71}$ .

---

*Example negative document*

---

Alice has 24 apples. In how many ways can she share them with Becky and Chris so that each of the three people has at least two apples?

(A) 105    (B) 114    (C) 190    (D) 210    (E) 380

Note: This solution uses the non-negative version for stars and bars. A solution using the positive version of stars is similar (first removing an apple from each person instead of 2).

This method uses the counting method of stars and bars (non-negative version). Since each person must have at least 2 apples, we can remove  $2 * 3$  apples from the total that need to be sorted. With the remaining 18 apples, we can use stars and bars to determine the number of possibilities. Assume there are 18 stars in a row, and 2 bars, which will be placed to separate the stars into groups of 3. In total, there are 18 spaces for stars + 2 spaces for bars, for a total of 20 spaces. We can now do  $\binom{20}{2}$ . This is because if we choose distinct 2 spots for the bars to be placed, each combo of 3 groups will be different, and all apples will add up to 18. We can also do this because the apples are indistinguishable.  $\binom{20}{2}$  is 190, therefore the answer is

**(C) 190**.

---

2214  
2215  
2216  
2217  
2218  
2219  
2220  
2221  
2222  
2223  
2224  
2225  
2226  
2227  
2228  
2229  
2230  
2231  
2232  
2233  
2234  
2235  
2236  
2237  
2238  
2239  
2240  
2241  
2242  
2243  
2244  
2245  
2246  
2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255  
2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267

Table 30: **TheoremQA Questions example.** Both the query and the positive document uses the Pigeonhole Principle to solve the problem. The negative document is retrieved by BM25, and it only requires simple arithmetics to solve.

<b>Query</b>	
	Mary is planning to bake exactly 10 cookies, and each cookie may be one of three different shapes – triangle, circle, and square. Mary wants the cookie shapes to be as diverse as possible. What is the smallest possible count for the most common shape across the ten cookies?
<b>Chain-of-thought reasoning to find documents</b>	
	This problem is similar to the formulation of Pigeonhole Principle. We can find solutions to other examples that can also be solved using this principle.
<b>Example positive document</b>	
	Arbitrarily place 19 points in a unit square and cover as many of these points as possible with a circle of diameter $\frac{\sqrt{2}}{3}$ . Question: At least how many points can be guaranteed to be covered? We can divide the unit square into 9 smaller squares, each with side length $1/3$ . Since there are 19 points and 9 smaller squares, by the Pigeonhole Principle, at least one of these smaller squares must contain at least 3 points. Now, consider a circle with diameter $\frac{\sqrt{2}}{3}$ . The diagonal of each smaller square has length $\frac{\sqrt{2}}{3}$ , so the circle can be placed inside the smaller square in such a way that it covers the entire square. Therefore, the circle can cover at least 3 points in that square. Therefore, the answer is 3.
<b>Example negative document</b>	
	Yanna baked twenty butter cookies and forty biscuits in the morning. In the afternoon, she baked ten butter cookies and twenty biscuits. How many more biscuits did she bake than butter cookies? She baked $20 + 10 = 30$ butter cookies. And, she baked $40 + 20 = 60$ biscuits. Therefore, she baked $60 - 30 = 30$ more biscuits than butter cookies. The answer is 30

2268 Table 31: **TheoremQA Theorems example.** The query uses the Pigeonhole Principle, which is the  
 2269 gold document from ProofWiki. The negative document is retrieved by BM25, which is not relevant  
 2270 to the question

2271	
2272	<i>Query</i>
2273	Mary is planning to bake exactly 10 cookies, and each cookie may be one of three different
2274	shapes – triangle, circle, and square. Mary wants the cookie shapes to be as diverse as possible.
2275	What is the smallest possible count for the most common shape across the ten cookies?
2276	<i>Chain-of-thought reasoning to find documents</i>
2277	This problem is similar to the formulation of Pigeonhole Principle. We may look into the
2278	principle and find hints about how this problem may be solved using that principle.
2279	
2280	<i>Example positive document</i>
2281	Pigeonhole Principle Tags: Pigeonhole Principle, Named Theorems, Combinatorics
2282	begin{theorem} Let $S$ be a finite set whose cardinality is $n$ . Let $S_1, S_2, \dots, S_k$ be a partition
2283	of $S$ into $k$ subsets. Then: at least one subset $S_i$ of $S$ contains at least $\lceil \frac{n}{k} \rceil$ elements where
2284	$\lceil \cdot \rceil$ denotes the ceiling function.
2285	end{theorem}
2286	begin{proof} AimForCont no subset $S_i$ of $S$ has as many as $\lceil \frac{n}{k} \rceil$ elements. Then the maximum
2287	number of elements of any $S_i$ would be $\lceil \frac{n}{k} \rceil - 1$ . So the total number of elements of $S$ would
2288	be no more than $k \left( \lceil \frac{n}{k} \rceil - 1 \right) = k \lceil \frac{n}{k} \rceil - k$ . There are two cases: $n$ is divisible by $k$ : $n$ is not
2289	divisible by $k$ . Suppose $k \nmid n$ . Then $\lceil \frac{n}{k} \rceil = \frac{n}{k} + 1$ is an integer and: $k \lceil \frac{n}{k} \rceil - k = n - k$ Thus:
2290	$ S  = \sum_{i=1}^k  S_i  \leq n - k < n$ This contradicts the fact that $ S  = n$ . Hence our assumption
2291	that no subset $S_i$ of $S$ has as many as $\lceil \frac{n}{k} \rceil$ elements was false. Next, suppose that $k \mid n$ .
2292	Then: $ S  = k \lceil \frac{n}{k} \rceil - k < \frac{k(n+k)}{k} - k = n$ and again this contradicts the fact that $ S  = n$ .
2293	In the same way, our assumption that no subset $S_i$ of $S$ has as many as $\lceil \frac{n}{k} \rceil$ elements was
2294	false. Hence, by Proof by Contradiction, there has to be at least $\lceil \frac{n}{k} \rceil$ elements in at least one
2295	$S_i \subseteq S$ . qed
2296	end{proof}
2297	
2298	<i>Example negative document</i>
2299	
2300	begin{definition}[Definition:Conditional/Semantics of Conditional]Let $p \implies q$ where $\implies$
2301	denotes the conditional operator. $p \implies q$ can be stated thus: * "If" $p$ is true "then" $q$ is
2302	true." * " $q$ is true "if" $p$ is true." * "(The truth of) $p$ "implies" (the truth of) $q$ ." * "(The
2303	truth of) $q$ "is implied by" (the truth of) $p$ ." * " $q$ "follows from" $p$ ." * " $p$ is true "only if" $q$ is
2304	true." * The latter one may need some explanation. $p$ can be either true or false, as can $q$ . But if
2305	$q$ is false, and $p \implies q$ , then $p$ can not be true. Therefore, $p$ can be true "only if" $q$ is also
2306	true, which leads us to our assertion. * " $p$ is true "therefore" $q$ is true." * " $p$ is true "entails"
2307	that $q$ is true." * " $q$ is true "because" $p$ is true." * " $p$ "may" be true "unless" $q$ is false." * "
2308	"Given that" $p$ is true, $q$ is true." * " $q$ is true "whenever" $p$ is true." * " $q$ is true "provided
2309	that" $p$ is true." * " $q$ is true "in case" $p$ is true." * " $q$ is true "assuming that" $p$ is true." * " $q$
2310	is true "on the condition that" $p$ is true." * Further colloquial interpretations can often be found
2311	in natural language whose meaning can be reduced down to " $p$ only if $q$ ", for example: * " $p$ is
2312	true "as long as" $q$ is true": Example::: "Mummy, can I go to the pictures?"::: "As long
2313	as" you've done your homework. Have you done your homework? No? Then you cannot go
2314	to the pictures."::: In other words::: "You can go to the pictures "only if" you have done
2315	your homework."::: Using the full language of logic::: "If it is true that you are going to the
2316	pictures, it is true that you must have done your homework." * " $p$ is true "as soon as" $q$ is
2317	true": "Are we going to this party, then?"::: "As soon as" I've finished putting on my
2318	makeup."::: The analysis is the same as for the above example of "as long as".
2319	end{definition}
2320	
2321	

2322  
2323  
2324  
2325  
2326  
2327  
2328  
2329  
2330  
2331  
2332  
2333  
2334  
2335  
2336  
2337  
2338  
2339  
2340  
2341  
2342  
2343  
2344  
2345  
2346  
2347  
2348  
2349  
2350  
2351  
2352  
2353  
2354  
2355  
2356  
2357  
2358  
2359  
2360  
2361  
2362  
2363  
2364  
2365  
2366  
2367  
2368  
2369  
2370  
2371  
2372  
2373  
2374  
2375

Table 32: **Example of Gemini reasoning on Sustainable Living**

---

**Query**

---

How good is it to reuse water from plant pots?  
I'm living in an apartment, and after I water my plants the water goes to plates below the pots. The pots are in a metallic structure above the plates, so I can take the plates to reuse the water (throwing it at the plants again).  
This reuse seems beneficial, because I think I can get rid of mosquitoes that would reproduce in the stagnated water. And also some nutrients of the soil (as well as earthworms) can return to the vase.  
Is there some negative points in doing that?

---

**LLM Reasoning steps**

---

## Plant Pot Water Reuse

\*\* (1) Essential Problem: \*\*

The user wants to know if it is beneficial or detrimental to reuse the water collected from plant pot drainage plates.  
They are concerned about potential mosquito breeding and the impact on plant growth.

\*\* (2) Reasoning: \*\*

\*\* Reused Water Composition: \*\*

- Potentially contains dissolved nutrients, minerals and salts from the soil.
- May contain microorganisms.
- Could harbor mosquito larvae if stagnant.

\*\* Negative Impact on Plant Growth: \*\*

- Stagnant water can attract mosquitoes.
- High salt concentrations from fertilizers can harm plants.
- Waterborne pathogens could potentially harm plants.

\*\* (3) Draft Answer: \*\*

## Reusing Water from Plant Pots: A Balancing Act

It's commendable that you're looking for ways to conserve water and potentially benefit your plants.  
Reusing water from plant pot drainage plates can be a good idea, but there are some drawbacks to consider:

- \* \*\*Mosquito Breeding:\*\* Stagnant water in the plates can become a breeding ground for mosquitoes, posing a health risk. Regularly empty and clean the plates to prevent this.
  - \* \*\*Salt Accumulation:\*\* Repeatedly using the same water can lead to a buildup of salts from fertilizers, which can harm plants.
  - \* \*\*Pathogen Risk:\*\* While less common, waterborne pathogens could potentially harm your plants. Using filtered or boiled water can minimize this risk.
-

2376  
2377  
2378  
2379  
2380  
2381  
2382  
2383  
2384  
2385  
2386  
2387  
2388  
2389  
2390  
2391  
2392  
2393  
2394  
2395  
2396  
2397  
2398  
2399  
2400  
2401  
2402  
2403  
2404  
2405  
2406  
2407  
2408  
2409  
2410  
2411  
2412  
2413  
2414  
2415  
2416  
2417  
2418  
2419  
2420  
2421  
2422  
2423  
2424  
2425  
2426  
2427  
2428  
2429

Table 33: Example of Gemini reasoning on Sustainable Living, continue

<i>Example positive document</i>	
	Soluble Salts
	Soluble salts may accumulate on the top of the soil, forming a yellow or white crust. A ring of salt deposits may form around the pot at the soil line or around the drainage hole. Salts may also build up on the outside of clay pots. In house plants, signs of excess soluble salts include reduced growth, brown leaf tips, dropping of lower leaves, small new growth, dead root tips, and wilting.
	Soluble salts are minerals dissolved in water. Fertilizer dissolved in water becomes a soluble salt. When water evaporates from the soil, the minerals or salts stay behind. As the salts in the soil become more and more concentrated, it becomes more difficult for plants to take up water. If salts build up to an extremely high level, water can be taken out of the root tips, causing them to die. High levels of soluble salts damage the roots directly, weakening the plant and making it more susceptible to attack from insects and diseases. One of the most common problems associated with high salt levels is root rot.
	The best way to prevent soluble salt injury is to stop the salts from building up. When watering, allow some water to drain through the container and then empty the saucer. Do not allow the pot to sit in water. If the drained water is absorbed by the soil, the salts that were washed out are reabsorbed through the drainage hole or directly through a clay pot.
	...

Table 34: The performance of retrieval models when using reasoning steps generated by GritLM as queries

	StackExchange						Coding		Theorem-based			Avg.	
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.		TheoT.
	<i>Sparse model</i>												
BM25 (Robertson et al., 2009)	33.1	38.5	19.2	28.1	16.7	16.2	20.7	19.7	13.3	3.3	12.1	8.7	19.1
	<i>Open-sourced models (&lt;1B)</i>												
BGE (Xiao et al., 2023)	17.7	29.8	18.7	22.9	12.6	11.6	17.6	23.5	3.2	4.2	15.4	11.3	15.7
Inst-L Su et al. (2022)	22.3	27.6	19.3	26.2	13.1	11.9	17.7	16.5	0.7	4.7	19.3	9.0	15.7
SBERT (Reimers & Gurevych, 2019)	17.0	22.0	15.2	23.9	9.7	9.3	13.1	24.2	1.7	3.8	15.4	9.4	13.7
	<i>Open-sourced models (&gt;1B)</i>												
E5 (Wang et al., 2023a)	24.3	36.5	18.5	21.2	10.6	14.9	15.0	27.0	0.3	4.0	20.0	17.4	17.5
SFR (Meng et al., 2024)	22.1	31.3	20.0	23.0	12.2	15.0	16.8	26.2	0.6	5.3	19.2	15.2	17.2
Inst-XL Su et al. (2022)	34.5	44.7	26.5	36.5	19.3	21.7	24.8	24.7	1.6	6.0	16.0	8.4	22.1
GritLM (Muennighoff et al., 2024)	24.6	31.6	18.8	21.0	14.2	10.9	21.5	28.9	5.0	5.5	18.6	17.0	18.1
Qwen (Li et al., 2023b)	25.7	35.8	19.5	29.7	12.1	18.1	22.0	22.9	7.8	2.6	17.9	21.8	19.7
	<i>Proprietary models</i>												
Cohere (Cohere)	21.8	26.4	22.2	21.5	15.1	15.8	17.1	21.3	1.9	4.7	14.4	9.9	16.0
OpenAI (OpenAI)	26.4	29.8	21.5	29.4	11.4	11.6	21.1	22.6	2.6	7.5	18.3	11.9	17.8
Voyage (Voyage AI)	28.4	29.0	19.9	28.3	9.6	17.7	17.2	29.2	1.7	7.2	23.1	10.6	18.5
Google (Lee et al., 2024)	25.6	34.4	22.1	30.0	14.8	19.3	16.5	28.9	1.6	8.5	16.9	12.5	19.3

Table 35: The performance of retrieval models when using reasoning steps generated by Gemini-pro as queries

	StackExchange							Coding		Theorem-based			Avg.
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.	TheoT.	
<i>Sparse model</i>													
BM25 (Robertson et al., 2009)	54.3	49.1	20.7	33.2	19.8	18.3	23.5	20.9	12.7	2.7	17.2	9.5	23.5
<i>Open-sourced models (&lt;1B)</i>													
BGE (Xiao et al., 2023)	28.4	34.7	21.8	25.1	13.5	13.4	19.1	24.7	5.3	4.0	17.7	13.8	18.4
Inst-L (Su et al., 2022)	35.9	41.3	22.8	34.0	17.8	16.7	16.4	17.4	0.9	6.1	23.2	12.1	20.4
SBERT (Reimers & Gurevych, 2019)	20.3	24.2	15.5	24.7	11.5	9.8	16.6	25.1	2.3	4.1	18.2	10.9	15.3
<i>Open-sourced models (&gt;1B)</i>													
E5 (Wang et al., 2023a)	29.7	42.0	18.1	19.2	13.8	17.1	14.1	27.6	0.2	5.9	20.5	23.9	19.3
SFR (Meng et al., 2024)	29.3	38.2	20.4	23.8	16.3	16.0	16.8	26.8	0.3	6.9	21.6	22.0	19.9
Inst-XL (Su et al., 2022)	44.5	48.7	26.8	34.7	22.3	22.0	26.5	24.8	1.5	7.0	17.4	11.6	24.0
GritLM (Muennighoff et al., 2024)	29.0	33.9	20.8	24.6	17.0	12.9	19.2	32.1	4.3	6.7	22.8	22.9	20.5
Qwen (Li et al., 2023b)	35.0	41.6	20.0	29.0	15.9	17.4	21.3	23.6	4.4	6.0	25.6	27.6	22.3
<i>Proprietary models</i>													
Cohere (Cohere)	30.2	33.2	22.9	27.1	16.6	19.3	21.6	20.7	3.4	5.2	18.2	15.6	19.5
OpenAI (OpenAI)	35.1	38.3	24.7	33.2	15.1	14.4	17.6	24.3	4.7	6.7	21.4	19.3	21.2
Voyage (Voyage AI)	37.5	40.0	22.8	29.7	13.6	17.9	19.6	32.3	2.3	5.3	26.7	17.6	22.1
Google (Lee et al., 2024)	35.6	38.4	23.6	33.1	16.6	22.5	16.2	32.3	2.6	8.0	23.8	17.4	22.5

Table 36: The performance of retrieval models when using reasoning steps generated by Llama3-70B as queries

	StackExchange							Coding		Theorem-based			Avg.
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.	TheoT.	
<i>Sparse model</i>													
BM25 (Robertson et al., 2009)	54.0	51.2	24.1	35.2	19.6	21.6	25.3	21.1	14.0	4.9	16.2	17.1	25.3
<i>Open-sourced models (&lt;1B)</i>													
BGE (Xiao et al., 2023)	28.8	38.6	22.7	23.5	15.0	16.5	22.2	25.4	7.5	5.7	21.7	16.7	20.3
Inst-L Su et al. (2022)	39.8	45.4	24.2	34.4	19.5	20.3	19.4	19.2	1.4	4.4	24.1	15.4	22.3
SBERT (Reimers & Gurevych, 2019)	20.2	25.5	16.9	24.1	10.3	11.7	16.6	24.7	6.6	3.8	19.7	13.8	16.1
<i>Open-sourced models (&gt;1B)</i>													
E5 (Wang et al., 2023a)	30.2	43.7	18.2	23.4	13.4	15.8	15.2	26.9	0.3	5.9	19.0	23.3	19.6
SFR (Meng et al., 2024)	26.1	38.4	21.4	25.3	14.7	15.8	17.7	27.1	0.5	6.9	20.3	21.5	19.7
Inst-XL Su et al. (2022)	45.2	51.2	32.1	39.8	21.1	26.2	28.0	25.3	1.7	6.8	19.5	13.1	25.8
GritLM (Muennighoff et al., 2024)	28.2	35.0	22.4	24.6	17.0	14.4	20.9	32.0	4.3	4.8	22.1	20.8	20.5
Qwen (Li et al., 2023b)	33.1	43.0	21.7	30.4	13.7	20.0	22.3	25.8	8.0	5.2	25.3	28.9	23.1
<i>Proprietary models</i>													
Cohere (Cohere)	34.8	35.1	26.4	27.6	17.9	22.4	24.0	24.5	4.2	6.6	19.7	19.7	21.9
OpenAI (OpenAI)	34.1	41.0	26.1	35.5	12.3	14.5	22.4	24.3	5.9	7.7	20.8	18.9	22.0
Voyage (Voyage AI)	37.1	41.1	23.8	26.9	13.6	18.8	24.2	31.9	1.4	6.9	26.0	21.4	22.8
Google (Lee et al., 2024)	39.5	44.0	26.4	37.0	18.7	25.0	17.5	31.1	3.1	8.8	21.2	21.4	24.5

Table 37: The performance of retrieval models when using reasoning steps generated by Claude-opus as queries

	StackExchange							Coding		Theorem-based			Avg.
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.	TheoT.	
<i>Sparse model</i>													
BM25 (Robertson et al., 2009)	54.2	51.7	23.2	38.3	22.4	20.3	25.8	20.0	19.8	4.0	18.5	17.6	26.3
<i>Open-sourced models (&lt;1B)</i>													
BGE (Xiao et al., 2023)	30.2	38.0	23.4	29.3	14.7	15.0	22.0	22.4	6.0	4.5	24.2	19.4	20.7
Inst-L Su et al. (2022)	37.8	40.0	24.9	34.2	19.6	17.6	18.9	17.8	1.5	4.0	24.5	18.2	21.6
SBERT (Reimers & Gurevych, 2019)	19.1	24.4	18.5	24.9	11.5	11.0	14.5	23.0	5.7	3.1	19.3	18.5	16.1
<i>Open-sourced models (&gt;1B)</i>													
E5 (Wang et al., 2023a)	31.6	43.3	20.9	25.8	13.6	17.4	14.6	29.1	0.3	6.8	23.1	26.4	21.1
SFR (Meng et al., 2024)	29.9	38.8	24.3	28.4	15.6	16.8	18.8	28.5	0.6	7.3	24.3	24.9	21.5
Inst-XL Su et al. (2022)	45.7	50.0	29.0	43.0	22.5	23.0	27.0	23.8	2.0	6.6	21.2	15.4	25.8
GritLM (Muennighoff et al., 2024)	32.4	36.4	22.5	28.8	19.5	15.8	24.1	31.1	6.3	5.7	25.0	26.5	22.8
Qwen (Li et al., 2023b)	35.8	45.1	22.0	34.4	15.8	19.5	26.4	24.6	5.7	6.2	27.0	30.9	24.5
<i>Proprietary models</i>													
Cohere (Cohere)	33.3	36.4	24.8	28.7	18.1	20.4	21.2	21.6	4.4	6.9	21.6	20.2	21.5
OpenAI (OpenAI)	36.5	39.6	25.6	36.7	13.2	14.8	22.6	25.5	5.4	7.7	22.5	21.2	22.6
Voyage (Voyage AI)	37.2	40.2	24.8	31.1	12.6	18.4	20.5	30.9	1.9	6.9	27.9	21.3	22.8
Google (Lee et al., 2024)	37.3	43.7	27.1	38.8	21.8	23.7	17.6	31.4	2.9	9.1	24.6	22.0	25.0



2484 Table 38: **The performance of retrieval models when using reasoning steps generated by GPT-4**  
 2485 **as queries**

	StackExchange							Coding		Theorem-based			Avg.
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.	TheoT.	
<i>Sparse model</i>													
BM25 (Robertson et al., 2009)	53.6	53.6	24.3	38.6	18.8	22.7	25.9	19.3	17.7	3.9	18.9	20.2	26.5
<i>Open-sourced models (&lt;1B)</i>													
BGE (Xiao et al., 2023)	29.3	41.8	23.4	27.8	11.7	14.4	21.6	24.2	8.2	7.5	24.3	25.1	21.6
Inst-L Su et al. (2022)	40.0	44.9	24.2	34.4	18.2	21.1	19.0	17.8	1.3	6.8	27.1	19.9	22.9
SBERT (Reimers & Gurevych, 2019)	18.9	26.0	17.5	27.1	9.1	9.9	17.3	24.3	10.2	5.0	21.4	22.9	17.5
<i>Open-sourced models (&gt;1B)</i>													
E5 (Wang et al., 2023a)	29.6	43.6	20.1	26.7	11.8	17.7	15.6	29.1	0.9	5.3	25.6	35.7	21.8
SFR (Meng et al., 2024)	26.2	39.1	21.5	28.3	13.4	16.8	19.5	28.4	1.5	7.1	25.9	33.2	21.7
Inst-XL Su et al. (2022)	47.3	51.4	29.7	40.5	20.5	25.3	27.0	24.4	2.1	8.4	21.8	16.4	26.2
GritLM (Muennighoff et al., 2024)	33.1	38.9	22.3	28.8	17.4	17.7	24.1	31.8	11.7	6.7	26.3	29.5	24.0
Qwen (Li et al., 2023b)	35.8	43.0	24.3	34.3	15.6	19.7	24.4	25.4	5.2	4.6	28.0	33.7	24.5
<i>Proprietary models</i>													
Cohere (Cohere)	31.8	36.6	25.6	29.3	16.1	22.0	23.2	20.8	7.0	7.0	23.0	25.2	22.3
OpenAI (OpenAI)	35.5	40.0	24.8	38.0	13.7	17.0	24.2	24.4	6.5	7.7	22.1	23.5	23.1
Voyage (Voyage AI)	36.9	42.7	24.5	34.3	13.9	21.2	22.1	31.4	2.1	6.6	29.0	27.4	24.4
Google (Lee et al., 2024)	36.5	45.3	25.6	38.1	19.1	25.0	18.0	31.1	3.6	10.1	27.4	29.6	25.8

2498 Table 39: **Long-context retrieval performance on unsplit web pages of StackExchange data.** The  
 2499 scores are reported in recall@1

	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Pony	Avg.
<i>Sparse models</i>									
BM25	10.7	15.4	10.7	8.4	7.4	22.2	10.7	5.4	11.4
<i>Open-sourced models (&lt;1B)</i>									
BGE	16.4	27.7	20.9	11.6	10.9	13.2	16.9	0.4	14.8
Inst-L	24.1	29.5	14.1	21.3	12.9	14.5	25.4	3.9	18.2
SBERT	25.6	34.1	18.9	15.8	10.9	15.0	18.0	1.2	17.4
<i>Open-sourced models (&gt;1B)</i>									
E5	29.9	36.3	26.2	46.7	17.3	14.5	32.2	1.1	25.5
SFR	30.3	37.0	24.3	47.7	17.3	14.5	35.0	2.0	26.0
Inst-XL	22.0	32.0	14.1	20.5	12.9	15.0	20.1	5.9	17.8
GritLM	37.5	39.4	25.7	35.3	17.8	20.1	32.4	0.0	26.0
Qwen	39.2	36.1	25.7	42.3	21.3	23.5	33.1	1.3	27.8
<i>Proprietary models</i>									
Cohere	31.5	35.5	18.0	20.5	9.9	15.8	15.2	0.8	18.4
OpenAI	33.1	32.3	24.8	35.1	11.9	11.5	26.3	0.0	21.9
Voyage	34.4	35.3	26.7	41.6	12.9	13.7	31.1	1.3	24.6
Google	30.9	38.0	21.8	30.7	12.9	19.2	25.6	0.3	22.4

2525 Table 40: **Statistics of StackExchange and Pony data before web pages and documentation are**  
 2526 **split.** For each dataset, we show the number of queries and documents, the average length of queries  
 2527 and documents, and the average number of gold documents.

	# Query	# Doc	Avg Query Len	Avg Doc Len	# Avg Gold Doc
Biology	103	524	115.2	9422.4	1.3
Earth Science	116	601	109.5	27312.3	1.6
Economics	103	516	181.5	11896.4	1.1
Psychology	101	512	149.6	12411.66	1.1
Robotics	101	508	818.9	14998.2	1.1
Stack Overflow	117	1858	478.3	40759.7	1.1
Sustainable Living	108	554	148.5	12077.7	1.2
Pony	112	577	102.6	1361.0	6.9

2538

2539

2540

Table 41: MiniLM (cross-encoder), Gemini and GPT-4 reranking scores based on BM25 top-10 or top-100 retrieval results.

2541

2542

2543

2544

2545

2546

2547

2548

2549

2550

2551

2552

2553

2554

2555

2556

Table 42: MiniLM (cross-encoder), Gemini and GPT-4 reranking scores based on Google retrieval top-10 or top-100 results.

2557

2558

2559

2560

2561

2562

2563

2564

2565

2566

2567

2568

2569

2570

2571

2572

Table 43: Gemini prompt to rerank documents.

2573

2574

2575

2576

2577

2578

2579

2580

2581

2582

2583

2584

2585

2586

2587

2588

2589

2590

2591

Gemini

The following passages are related to the query: {query}

[1]. {doc 1}

[2]. {doc 2}

...

First identify the essential problem in the query.

Think step by step to reason about why each document is relevant or irrelevant.

Rank these passages based on their relevance to the query.

Please output the ranking result of the most relevant {k} passages as a list,

where the first element is the id of the most relevant passage,

the second element is the id of the second most element, etc.

Please strictly follow the format to output a list of {k} ids:

"""

[...]

"""

2592  
2593  
2594  
2595  
2596  
2597  
2598  
2599  
2600  
2601  
2602  
2603  
2604  
2605  
2606  
2607  
2608  
2609  
2610  
2611  
2612  
2613  
2614  
2615  
2616  
2617  
2618  
2619  
2620  
2621  
2622  
2623  
2624  
2625  
2626  
2627  
2628  
2629  
2630  
2631  
2632  
2633  
2634  
2635  
2636  
2637  
2638  
2639  
2640  
2641  
2642  
2643  
2644  
2645

Table 44: **Results of retrieval models copied from MTEB Muennighoff et al. (2023) for easier reference.** Argu. refers to ArguAna, Climate. refers to ClimateFEVER, CQA. refers to CQADup-stackRetrieval, FIQA. refers to FIQA2018, Hot. refers to HotpotQA, MS. refers to MSMARCO, NF. refers to NFCorpus, Quora refers to QuoraRetrieval, SCI. refers to SCIDOCS, Sci. refers to SciFact, Touche. refers to Touche2020, TREC. refers to TRECCOVID. Except BM25, whose results are from Thakur et al. (2021), all other results are from Muennighoff et al. (2023).

	Argu.	Climate.	CQA.	DBPedia	FEVER	FIQA.	Hot.	MS.	NF.	NQ	Quora.	SCI.	Sci.	Touche.	TREC.	Avg.
<i>Sparse model</i>																
BM25 (Robertson et al., 2009)	31.5	21.3	29.9	31.3	75.3	23.6	60.3	22.8	32.5	32.9	78.9	15.8	66.5	36.7	65.6	41.6
<i>Open-sourced models (&lt;1B)</i>																
BGE (Xiao et al., 2023)	63.5	36.6	42.2	44.1	87.2	45.0	74.1	42.6	38.1	55.0	89.1	22.6	74.6	24.8	74.8	54.3
Inst-L Su et al. (2022)	57.1	27.7	43.8	36.7	72.7	45.5	55.2	39.7	34.1	50.1	88.4	18.6	64.3	21.6	58.1	47.6
SBERT (Reimers & Gurevych, 2019)	46.5	22.0	45.0	32.1	50.9	50.0	39.3	39.8	33.3	50.5	87.5	23.8	65.6	19.9	51.3	43.8
<i>Open-sourced models (&gt;1B)</i>																
E5 (Wang et al., 2023a)	61.9	38.4	43.0	48.9	87.8	56.6	75.7	43.1	38.6	63.5	89.6	16.3	76.4	26.4	87.3	56.9
SFR (Meng et al., 2024)	67.2	36.4	46.5	49.1	89.4	60.4	77.0	43.4	41.9	69.9	89.8	19.9	77.7	29.0	87.6	59.0
Inst-XL Su et al. (2022)	55.7	26.5	43.1	40.2	70.0	47.0	55.9	41.6	36.0	57.2	88.9	17.4	64.6	23.4	71.4	49.3
GritLM (Muennighoff et al., 2024)	63.2	30.9	49.4	46.6	82.7	60.0	79.4	42.0	40.9	70.3	89.5	24.4	79.2	27.9	74.8	57.4
Qwen (Li et al., 2023b)	62.7	44.0	40.6	48.0	93.4	55.3	72.3	41.7	38.3	61.8	89.6	27.7	75.3	20.3	72.7	56.2
<i>Proprietary models</i>																
Cohere (Cohere)	61.5	38.4	41.5	43.4	89.0	42.2	70.7	42.9	38.6	61.6	88.7	20.3	71.8	32.4	81.9	55.0
Voyage (Voyage AI)	64.1	32.7	46.6	46.0	91.5	59.8	70.9	40.6	40.3	65.9	87.4	24.3	80.0	39.2	85.1	58.3
OpenAI (OpenAI)	58.1	30.3	47.5	44.8	87.9	55.0	71.6	40.2	42.1	61.3	89.1	23.1	77.8	23.4	79.6	55.4
Google (Lee et al., 2024)	62.2	33.2	48.9	47.1	87.0	59.2	71.3	32.6	40.3	61.3	88.2	20.3	75.4	25.9	82.6	55.7

Table 45: QA inference prompt

---

Problem:  
{problem\_description}

Documents:  
{retrieved\_doc}

Based on the provided documents, write an answer to the problem.

---

2646  
2647  
2648  
2649  
2650  
2651  
2652  
2653  
2654  
2655  
2656  
2657  
2658  
2659  
2660  
2661  
2662  
2663  
2664  
2665  
2666  
2667  
2668  
2669  
2670  
2671  
2672  
2673  
2674  
2675  
2676  
2677  
2678  
2679  
2680  
2681  
2682  
2683  
2684  
2685  
2686  
2687  
2688  
2689  
2690  
2691  
2692  
2693  
2694  
2695  
2696  
2697  
2698  
2699

Table 46: QA evaluation prompt

---

\_\_\_\_\_ PROBLEM START \_\_\_\_\_  
 {predicted\_answer}  
 \_\_\_\_\_ PROBLEM END \_\_\_\_\_  
 \_\_\_\_\_ STUDENT ANSWER START \_\_\_\_\_  
 {predicted\_answer}  
 \_\_\_\_\_ STUDENT ANSWER END \_\_\_\_\_  
 \_\_\_\_\_ REFERENCE ANSWER START \_\_\_\_\_  
 {gold\_answer}  
 \_\_\_\_\_ REFERENCE ANSWER END \_\_\_\_\_

Criteria:

0 - The student’s answer is completely irrelevant or blank.  
 10 - The student’s answer addresses about 10% of the reference content.  
 20 - The student’s answer addresses about 20% of the reference content.  
 30 - The student’s answer addresses about 30% of the reference content.  
 40 - The student’s answer addresses about 40% of the reference content.  
 50 - The student’s answer addresses about 50% of the reference content.  
 60 - The student’s answer addresses about 60% of the reference content.  
 70 - The student’s answer addresses about 70% of the reference content.  
 80 - The student’s answer addresses about 80% of the reference content.  
 90 - The student’s answer addresses about 90% of the reference content.  
 100 - The student’s answer addresses about 100% of the reference content.

Use the following format to give a score:  
 REASON:  
 Describe why you give a specific score  
 SCORE:  
 The score you give, e.g., 60  
 Do not say anything after the score.

---

Table 47: **Question answering results with different retrievers.** We use the GPT-4 as the generation models, and evaluate the answers with Claude-3.5-sonnet. Results show that stronger retrieval typically result in better QA results, indicating the helpfulness of the annotated documents for addressing the posts in StackExchange.

Retriever	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Average
None	62.6	64.9	57.8	58.2	53.5	64.4	49.8	58.7
BM25	64.3	65.2	57.8	54.6	60.2	65.9	56.5	60.6
SBERT	62.2	63.2	59.7	65.6	55.7	59.9	<b>61.0</b>	61.0
Qwen	<b>69.4</b>	<b>68.3</b>	<b>62.0</b>	<b>66.3</b>	<b>64.2</b>	<b>69.5</b>	55.5	<b>65.0</b>
Oracle	80.0	81.2	73.6	77.8	70.5	80.4	70.6	76.3

Table 48: Downstream results of gpt-4o-2024-08-06 on TheoremQA Questions, TheoremQA Theorems, and AoPS. The results are averaged across 5 random seeds and we report the standard deviation in the subscript.

Model	Setting	TheoQ.	TheoT.	AoPS
GPT-4o	None	76.3 <sub>1.6</sub>	82.1 <sub>3.0</sub>	36.6 <sub>2.6</sub>
	Random	76.5 <sub>1.3</sub>	85.0 <sub>1.8</sub>	37.4 <sub>2.6</sub>
	Qwen	76.4 <sub>2.2</sub>	82.6 <sub>2.2</sub>	36.1 <sub>1.9</sub>
	Oracle	79.3 <sub>0.4</sub>	89.7 <sub>1.7</sub>	37.2 <sub>3.8</sub>

2700  
2701  
2702  
2703  
2704  
2705  
2706  
2707  
2708  
2709  
2710  
2711  
2712  
2713  
2714  
2715  
2716  
2717  
2718  
2719  
2720  
2721  
2722  
2723  
2724  
2725  
2726  
2727  
2728  
2729  
2730  
2731  
2732  
2733  
2734  
2735  
2736  
2737  
2738  
2739  
2740  
2741  
2742  
2743  
2744  
2745  
2746  
2747  
2748  
2749  
2750  
2751  
2752  
2753

Table 49: The performance of retrieval models on **BRIGHT** measured by Precision@10.

	StackExchange						Coding		Theorem-based			Avg.	
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.		TheoT.
<i>Sparse model</i>													
BM25	7.6	12.4	7.1	6.0	5.6	8.0	6.1	6.0	7.9	3.1	2.2	1.3	6.1
<i>Open-sourced models (&lt;1B)</i>													
BGE	5.9	8.6	8.0	8.2	5.1	4.9	5.7	6.3	5.9	3.2	2.7	1.9	5.5
Inst-L	6.8	8.5	7.6	9.4	4.5	5.7	6.3	5.4	1.2	4.0	4.4	2.6	5.5
SBERT	7.2	7.4	8.3	10.1	4.5	5.0	6.8	6.5	7.0	3.2	3.8	2.7	6.0
<i>Open-sourced models (&gt;1B)</i>													
E5	8.9	10.2	8.1	8.6	7.2	4.6	6.9	6.9	4.9	4.2	5.7	6.5	6.9
SFR	9.3	10.3	9.0	10.3	7.1	5.9	7.7	6.6	2.1	4.2	5.2	6.5	7.0
Inst-XL	10.0	13.8	10.6	11.0	6.7	8.8	9.2	6.3	4.6	4.7	3.3	1.9	7.6
GritLM	11.1	12.7	9.2	10.8	6.8	6.0	6.8	7.5	17.9	4.6	5.7	5.3	8.7
Qwen	13.5	14.1	8.2	11.2	5.8	10.1	6.1	6.3	9.7	7.1	6.2	7.3	8.8
<i>Proprietary models</i>													
Cohere	8.6	10.3	10.4	10.2	6.5	7.8	8.5	6.5	1.8	3.4	3.0	1.7	6.6
OpenAI	11.4	11.0	9.8	12.4	5.8	6.3	9.0	6.3	2.4	4.5	5.1	2.9	7.2
Voyage	11.0	9.9	9.6	11.0	5.6	7.3	7.5	7.5	1.1	5.0	5.6	3.2	7.0
Google	10.3	12.2	8.9	11.4	5.6	8.3	7.9	6.9	3.6	5.0	4.8	4.2	7.4

Table 50: The performance of retrieval models on **BRIGHT** measured by Recall@10.

	StackExchange						Coding		Theorem-based			Avg.	
	Bio.	Earth.	Econ.	Psy.	Rob.	Stack.	Sus.	Leet.	Pony	AoPS	TheoQ.		TheoT.
<i>Sparse model</i>													
BM25	21.8	31.4	16.8	15.5	19.4	16.8	21.1	29.5	3.6	6.0	11.4	9.0	16.9
<i>Open-sourced models (&lt;1B)</i>													
BGE	15.1	27.0	16.2	18.4	14.4	12.1	17.0	31.3	3.0	7.2	14.6	11.0	15.6
Inst-L	18.8	27.8	17.5	26.8	15.6	15.1	17.4	23.6	0.7	8.2	22.8	14.8	17.4
SBERT	18.1	25.4	18.7	23.9	11.0	12.7	18.8	31.4	3.5	5.8	20.8	15.7	17.2
<i>Open-sourced models (&gt;1B)</i>													
E5	22.0	29.4	18.4	18.3	18.7	11.9	23.0	34.6	2.4	8.2	27.2	34.8	20.7
SFR	22.7	30.6	21.7	25.3	19.8	16.0	25.3	32.9	1.1	7.7	25.1	35.6	22.0
Inst-XL	27.3	38.0	25.4	35.6	22.0	21.1	23.9	31.8	2.5	8.9	16.6	9.8	21.9
GritLM	30.3	38.8	18.3	26.9	21.3	15.1	23.4	36.3	8.2	9.4	26.2	26.6	23.4
Qwen	38.2	40.6	18.5	29.5	14.5	22.4	17.4	32.1	4.6	14.8	30.0	39.4	25.2
<i>Proprietary models</i>													
Cohere	23.1	29.6	22.4	25.1	17.7	21.2	23.4	31.1	0.9	7.1	15.4	9.3	18.9
OpenAI	29.1	30.5	24.5	36.0	16.2	15.8	25.1	29.4	1.3	8.1	25.8	17.1	21.6
Voyage	29.3	31.2	21.0	31.0	15.0	17.5	20.5	41.5	0.6	8.7	28.5	15.4	21.7
Google	26.1	36.9	20.6	31.4	17.7	21.6	23.7	33.5	1.9	10.4	24.0	22.1	22.5