

Automatic Inter-document Multi-hop Scientific QA Generation

Seungmin Lee^{1,2}, Dongha Kim^{1,2}, Yuni Jeon^{1,2}, Junyoung Koh¹, and Min Song^{1,2*}

¹Yonsei University, Seoul, Republic of Korea

²OnomaAI, Seoul, Republic of Korea

{elplaguister, eastha, jeonuni, solbon1212, min.song}@yonsei.ac.kr

Abstract

Existing automatic scientific question generation studies mainly focus on single-document factoid QA, overlooking the inter-document reasoning crucial for scientific understanding. We present AIM-SciQA, an automated framework for generating multi-document, multi-hop scientific QA datasets. AIM-SciQA extracts single-hop QAs using large language models (LLMs) with machine reading comprehension and constructs cross-document relations based on embedding-based semantic alignment while selectively leveraging citation information. Applied to 8,211 PubMed Central papers, it produced 411,409 single-hop and 13,672 multi-hop QAs, forming the IM-SciQA dataset. Human and automatic validation confirmed high factual consistency, and experimental results demonstrate that IM-SciQA effectively differentiates reasoning capabilities across retrieval and QA stages, providing a realistic and interpretable benchmark for retrieval-augmented scientific reasoning. We further extend this framework to construct CIM-SciQA, a citation-guided variant achieving comparable performance to the Oracle setting, reinforcing the dataset’s validity and generality.

Keywords: Scientific Question Answering, Automatic Question Generation, Multi-hop QA, Dataset Construction

1. Introduction

Large language models (LLMs) play a pivotal role in scientific research, significantly advancing from simple text completion tools to sophisticated systems capable of complex reasoning and autonomous problem-solving. This progress has enabled "co-scientist" to independently generate hypotheses, conduct experiments, and write papers (Lu et al., 2024; Gottweis et al., 2025). LLMs now achieve near-human performance on question-answering (QA) benchmarks, and retrieval-augmented generation (RAG) facilitates sophisticated multi-document reasoning (Lewis et al., 2020).

Despite these advances, multi-hop scientific QA remains a crucial challenge for enabling machines to reason across interconnected documents (Liu et al., 2023; Wadden et al., 2022; Shi et al., 2024; Dasigi et al., 2021; Park et al., 2025). This difficulty arises from the vast volume of literature, domain-specific terminology, and fragmented information distribution, all of which make cross-document reasoning inherently complex. Moreover, existing datasets are scarce and limited in scope: most benchmarks focus on single-document QA, failing to reflect the complexity of real scientific inquiries that span multiple sources. Although recent datasets such as M3SciQA (Li et al., 2024) have begun addressing multi-document scenarios, the landscape remains sparse, particularly for multi-hop QA where answers must be composed from multiple pieces of evidence. This data scarcity continues to hinder the development and evaluation of

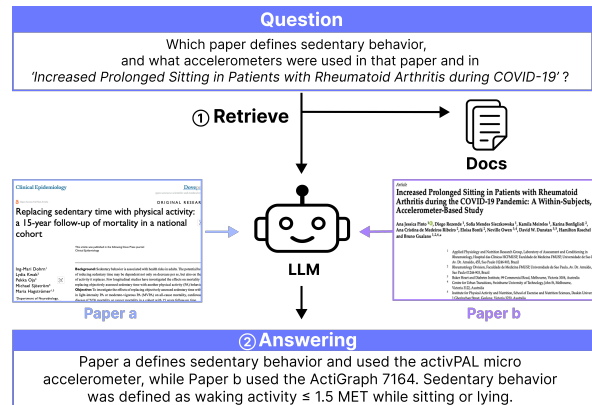


Figure 1: Example of Inter-document Multi-hop QA: ① the model retrieves relevant documents and ② integrates complementary evidence to answer.

robust scientific QA models.

A key bottleneck in scientific QA is curating datasets that are both diverse and challenging. Manual annotation of scientific documents is costly and requires domain experts. For instance, PubMedQA (Jin et al., 2019) combined 1,000 expert-labeled items with over 210k automatically generated instances, underscoring the scalability of automatic question generation (QG). However, most automated QG efforts focus on general-domain sources such as Wikipedia, with little attention to primary scientific documents. Scientific documents present additional challenges, including specialized terminology, dense exposition, explicit citation networks, and implicit relations across papers. Effective QG must explicitly leverage or robustly accommodate these characteristics.

*Corresponding author

Table 1: Comparison with other datasets. **M hop**: Multi Hop, **M Doc**: Multi Documents, **Q Gen**: Question Generation method, **Free-form**: answers are open-ended.

Dataset	M Hop	M Doc	#QA	QA Gen	Free-form	Domain	Source
PubMedQA	X	X	211k	Human	X	Biomedical	PubMed
SciQAG-24D	X	X	180k	LLM	○	Scientific Literature	Scientific Literature
WikiHop	○	○	51k	Rule-based	X	General/Wiki	Wiki
HotpotQA	○	○	113k	Human	○	General/Wiki	Wiki
DocHop-QA	○	○	11k	LLM	X	Biomedical/Scientific	PubMed
M3SciQA	○	○	1.4k	Human	X	NLP Papers	NLP Papers (+cited)
IM-SciQA (ours)	○	○	13.4K	LLM	○	Biomedical	PubMed

In this paper, we propose Automatic Inter-document Multi-hop Scientific QA generation (AIM-SciQA), an automated QG methodology tailored for scientific QA. Our approach uses a specialized LLM in a machine reading comprehension (MRC) paradigm to extract potential QA pairs from individual papers, then performs cross-paper matching to construct multi-document questions by aggregating information across documents. **AIM-SciQA** considers both explicit citation links and implicit connections, enabling automatic generation of complex multi-document multi-hop QA. We apply this pipeline to the biomedical domain, generating approximately 411,409, single-hop QA pairs from individual papers and deriving tens of thousands of multi-document, multi-hop examples.

To our knowledge, this is the first large-scale scientific QA dataset with automatically generated multi-hop questions spanning multiple documents in a specialized domain. Extensive experiments on the IM-SciQA and CIM-SciQA datasets demonstrate that AIM-SciQA produces high-quality multi-hop QA pairs and provides a realistic benchmark for evaluating retrieval and reasoning performance of large language models. With AIM-SciQA, we aim to advance the state of Scientific QA by enabling efficient dataset creation and providing a new benchmark to drive research in complex multi-document question answering on scientific texts. The datasets **IM-SciQA** and **CIM-SciQA** are publicly available on Hugging Face.

In summary, our contributions are as follows:

- **AIM-SciQA Framework**: We introduce a novel automated QA Generation pipeline that generates inter-document, multi-hop QA dataset from scientific literature.
- **IM-SciQA Dataset**: By leveraging AIM-SciQA framework, we automatically created IM-SciQA, large-scale scientific inter-document multi-hop QA Dataset including 13,672 QAs from PubMed Central.
- **Comprehensive Analyses**: We conducted comprehensive evaluations of recent Encoders and LLMs on IM-SciQA, revealing the

effectiveness of our restricted retrieval setting and completeness of inter-document QA.

2. Related Works

2.1. Scientific QA Datasets

Early scientific QA centered on single-document understanding as seen in PubMedQA (Jin et al., 2019) (abstract-level, yes/no/maybe), COVID-QA (Möller et al., 2020) (expert-curated, topic-specific), ScienceQA (Lu et al., 2022) (multimodal, MCQ), and QASPER (Dasigi et al., 2021) (full-paper comprehension). Domain-focused resources such as QASA (Lee et al., 2023), SPIQA (Pramanick et al., 2024) and BioASQ (Kriithara et al., 2023) expanded coverage but largely retained single-document or constrained formats. Multi-hop and cross-document settings emerged with M3SciQA (Li et al., 2024), and domain extensions (e.g., ContractNLI (Koreeda and Manning, 2021), ChemRxivQuest (Amiri and Bocklitz, 2025)). Recent biomedical efforts—DocHop-QA (Park et al., 2025) and PeerQA (Baumgärtner et al., 2025) advance realism and scale but often rely on predefined templates or limited question types.

Overall, existing benchmarks as shown in Table 1, either confine reasoning to a single document or constrain question diversity, motivating free-form, multi-document, multi-hop QA benchmark.

2.2. Automatic Question Generation

To reduce annotation cost and increase realism, automated QG has gained traction. LIQUID (Lee et al., 2023) scales list-style QG on wikipedia but not multi-hop. SciQAG (Wan et al., 2024) shows LLM-based QG at scale across scientific domains while remaining primarily single-document. Multi-hop QG without rigid templates has been explored via end-to-end question rewriting (E2EQR) (Hwang et al., 2024) and type-aware CoT with semantic extraction (TASE-CoT) (Lin et al., 2024). DocHop-QA applies LLMs to biomedical multi-document QG yet guides generation with 11 predefined concepts, potentially limiting naturalness.

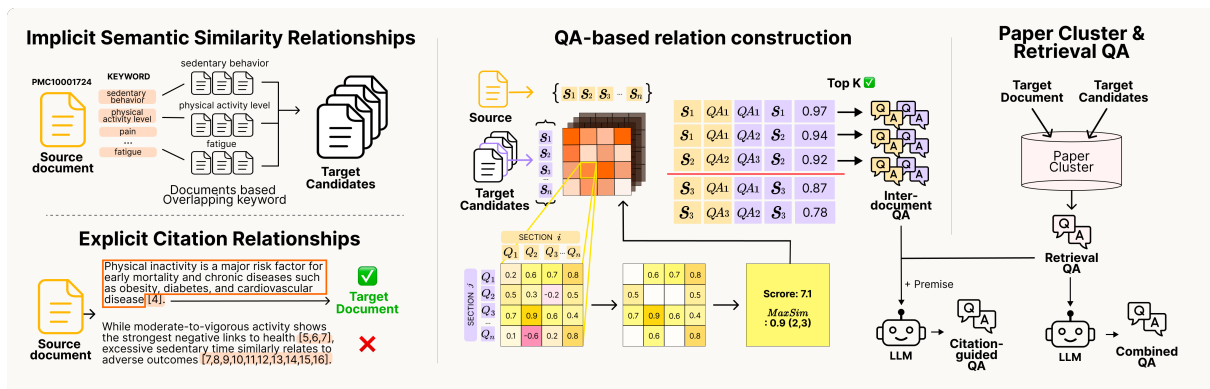


Figure 2: Illustration of the relation construction process in AIM-SciQA. For IM-SciQA, candidate relations are created from papers sharing overlapping keywords with the source paper (**implicit semantic similarity relationships**), while CIM-SciQA constructs relations directly from explicit citation links (**explicit citation relationships**). From these candidate relations, **QA-based relation construction** (Section 4.3) process selects promising single-hop QA (made by Section 4.2) pairs across papers to form inter-document QA candidates, and a **paper cluster** is built to define the target paper and its corresponding **retrieval QA**.

Our framework focuses on automatically constructing inter-document, multi-hop QA for rigorous evaluation of inter-document synthesis.

3. Problem Statement

Inter-document Multi-hop QA (IM-QA) aims to answer a question by integrating information distributed across multiple documents, as illustrated in Figure 1. Let D denote the collection of all documents, Q the set of questions, and A the set of possible answers. For a given question $q \in Q$, the goal is to identify a subset of documents $D_q \subseteq D$ and to derive an answer $a \in A$ from them. Formally, this task can be expressed as follows:

$$a = f(q, D_q), \quad \text{where } |D_q| > 1. \quad (1)$$

f denotes the reasoning function that derives an answer a by integrating evidence from multiple relevant documents D_q . This process can be divided into two stages—(1) **retrieving** a subset of relevant documents D_q from the entire collection D and (2) **synthesizing** the aggregated information within D_q to produce the final answer. We adopt a source-provided IM-QA setting, where each question is associated with a predefined set of source documents S_q . This constraint enables controlled and interpretable retrieval evaluation while retaining the core aspect of multi-document reasoning.

4. AIM-SciQA

We propose **AIM-SciQA**, an automated framework for constructing IM-QA over scientific documents. We additionally build **IM-SciQA**, a biomedical paper-based IM-QA dataset. The following section

describes how IM-SciQA is constructed and concurrently provides a detailed account of the AIM-SciQA. Our framework comprises as follows: (1) paper selection, (2) **single-hop QA (SHQA) generation**, (3) **document relation construction**, and (4) **multi-hop QA (MHQA) generation**.

4.1. Paper Selection

We collected 20,281 biomedical research articles available in PubMed Central (PMC) ¹ between 2019 and 2024 to construct the initial corpus. We then applied filtering criteria to ensure that each paper provides complete metadata, including full text, abstract, and reference information. Because cited papers are not always available within the PMC subset, we restricted relation construction to citations pointing to papers that are also accessible in PMC and retained only papers with at least three such accessible citations. After applying these filters, 8,211 papers were retained as valid source documents for dataset construction.

4.2. Single Hop QA Generation

Then, we automatically generated single-hop QA pairs that can be answered within a single paper. This step provides the atomic units for constructing multi-hop QA by enabling effective combination of information across papers. To identify important information in each paper and extract it in a QA format, we utilized *Qwen2.5-7B-Instruct* (Team, 2024). Specifically, given a document D and one of its sections S , an instruction-tuned LLM_θ is prompted to produce a question-answer-evidence triplet:

$$LLM_\theta(D, S) \longrightarrow (Q, A, E) \quad (2)$$

¹ <https://pmc.ncbi.nlm.nih.gov/>




Figure 3: Average Single-hop QA Count by Section




Figure 4: Total Single-hop QA Count by Section

Table 2: Statistics of the Single-hop QA

Item	Value
Count	411,409
Avg. Question Length	88.75
Avg. Answer Length	126.92
Avg. QA per Paper	55.67
Avg. QA per Section	10.49

where Q asks for an information slot, A provides its value, and E is the supporting sentence in the text.

We designed a prompt to generate the list of triplets from each paragraph, where the evidence is the original sentence that the model referenced when creating the answer. The prompt used for SHQA generation is provided in Appendix A.1. To ensure quality, we only retained QA pairs where the evidence sentence exactly matched the paper’s text, eliminating hallucinations. This process yielded 457K SHQA pairs; their statistics are summarized in Figure 3, Figure 4, and Table 2.

Quality Assurance Quality verification was conducted in two complementary phases. First, a manual evaluation was performed by four graduate students on 210 randomly sampled QA pairs, assessing appropriateness on a 0–0.5–1 scale. Among them, 202 pairs were rated perfect, achieving 0.94 Gwet’s AC1 score, which corresponds to almost perfect agreement levels (Gwet, 2008; Landis and Koch, 1977). Closer inspection of the remaining error cases revealed that most errors occurred when the questions and evidence were not semantically aligned. Notably, these mismatched cases showed an average $\text{Sim}(Q, A) - \text{Sim}(Q, E)$ approximately 0.1 lower than that of well-matched pairs. Based on this observation, we implemented an automatic filtering procedure that removed the lowest 10% of QA pairs by this similarity difference, effectively discarding samples most likely to contain misaligned evidence–answer mappings. After this two-stage verification, a total of 411,409 SHQA pairs were retained as the final SHQA dataset.

4.3. Document Relation Construction

We constructed inter-document relations for multi-hop QA through a unified pipeline primarily based on **QA-based semantic matching**, with an optional **Citation-guidance** as depicted in Figure 2. Each document is represented as a collection of SHQA triplets (Q, A, E) , which act as proxies for section-level information units. These triplets are encoded at the Q , A , and QA levels for our PMC-derived SHQA corpus using the *MedEmbed-large-v0.1* encoder (Balachandran, 2024), a model trained on the PMC-based BioASQ-QA dataset. We represent each section by its question embeddings, not QA or answer embeddings. This choice is motivated by our goal: not to find sections with identical facts but to locate those whose information property or semantic attribute required by the question is similar. In other words, question embeddings enable us to align sections based on the similarity of their underlying informational attributes, rather than overly trivial content.

Formally, each document D_i is defined as the set of its sections:

$$D_i = \{S_1^{(i)}, S_2^{(i)}, \dots, S_x^{(i)}\}. \quad (3)$$

Each section contains multiple embedded Question units. For sections $S_i^{(1)}$ and $S_j^{(2)}$, let $\mathbf{u}_{i,p}$ denote the embedding of the p -th Q in $S_i^{(1)}$ and $\mathbf{v}_{j,q}$ that of the q -th Q in $S_j^{(2)}$. We define the pairwise similarity matrix as:

$$K_{i,j}[p, q] = \text{CosSim}(\mathbf{u}_{i,p}, \mathbf{v}_{j,q}). \quad (4)$$

The section-level similarity is then computed as a weighted sum over pairs whose similarity exceeds the threshold $\tau = 0.3$:

$$\text{Sim}(S_i^{(1)}, S_j^{(2)}) = \sum_{p=1}^{n_i} \sum_{q=1}^{m_j} \mathbf{1}(K_{i,j}[p, q] \geq \tau) \cdot |K_{i,j}[p, q]|. \quad (5)$$

This threshold suppresses irrelevant matches while preserving meaningful QA-level alignments and can be adjusted depending on the dataset characteristics or the embedding model used. For each

document pair (D_1, D_2) , we compute Sim across all section pairs and select the **Top- K section pairs** with the highest similarity scores. For each selected section pair, we identify the QA pair with the highest similarity score and use this score to represent the strength of the relation between the corresponding sections of the two documents.

Finally, we enforce diversity via top-k selection on mapped scores, keeping at most three IM-QA instances per document in the final set. As a result, we collected 14,365 candidate IM-QA items.

Paper Cluster Construction For each source D_1 , we form a compact multiple-choice cluster of exactly 30 documents comprising the target document and up to 29 distractors. We start from the **keyword-overlap list** that consists of papers sharing keywords with D_1 , capped at 29 and, if shorter, pad with random non-candidates without duplication. This cluster serves two purposes: (i) defining the answer set for first-hop target identification, and (ii) providing a contrast set for selecting a target-side retrieval QA whose content is most distinctive relative to the cluster.

Retrieval QA Selection To select a retrieval QA that uniquely identifies the target paper within the paper cluster, we aggregate all QA units of the target paper together with all QA units from the paper cluster. Then, we compute QA embeddings for each item and calculate pairwise cosine distances across the entire pool. For each QA in the target paper, we compute its distinctiveness by aggregating cosine distances to all QAs from other papers in the cluster, and select the QA with the highest total distance. This QA is designated as the retrieval QA, as it remains grounded in the target paper while being most distinct from semantically similar candidates.

Citation-guided Relation Construction The citation-guided relation construction process follows the same procedure as above but replaces the target paper selection step and paper cluster construction step with citation evidence and cited paper pool, respectively. For target paper selection step, instead of forming candidates through keyword overlaps, we enumerate all in-text citations in the source document D_1 and treat sentences containing a *single citation marker* as grounded links to their referenced document D_2 . Formally, if \mathcal{V} denotes the document set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents directed citation edges, each sentence s^* in D_1 containing one reference marker deterministically maps to its bibliographic target D_2 . The paper cluster construction step is also replaced to the list of all cited paper that available publicly. Thus,

Table 3: Statistics of the Multi-hop QA

Item	Dev	Test
Count	13,372	300
Avg. Retrieval Q Length	15.05	15.16
Avg. Inter-doc Q Length	28.68	28.65
Avg. Combined Q Length	48.52	49.21
Avg. Inter-doc A Length	28.13	29.98
Avg. Combined A Length	43.00	45.33

the citation-based variant retains the same embedding alignment and selection process as the QA-based approach, while substituting the candidate generation mechanism with citation-derived pairs. To simulate and validate this procedure, we additionally developed Citation-guided Inter-document Multi-hop Scientific QA dataset (**CIM-SciQA**) with 74 items, built upon curated 51 papers which have all cited papers within our collected 8211 papers.

4.4. Multi-hop Question Generation

AIM-SciQA finally constructs multi-document, multi-hop questions from section-anchored (Q, A, E) tuples in the source and target document. Instead of simply splicing two QAs, we construct an explicit multi-step reasoning chain that integrates knowledge across papers. Prior to generation, a pre-validation filters out non-relational pairs (e.g., lacking comparison, causation, or inference), resulting in 13,672 valid instances out of 14,365 candidates with 693 pre-rejected. The remaining pairs then undergo a stepwise process to produce the final IM-QA. Table 3 summarizes the statistics of the resulting multi-hop QA dataset. The full prompting template for question generation is provided in Appendix A.2.

Step 1. Find Target Paper We first reconstruct the SHQA associated with the target paper into a paper identification question, whose unique answer is the target document itself. Consequently, the answer must include information that is uniquely attributable to the target paper within the given retrieval scope. To achieve this, we utilize the Retrieval QA defined during the paper cluster construction stage. Specifically, the Retrieval QA is reformulated so that its question explicitly highlights the target paper’s distinctive topic, variables, or findings, as in the following example: *“Which paper measured the incidence of pneumonia among Japanese participants in 2023?”*

Step 2. Generate Inter-document QA Next, we construct a comparative or integrative question that

explicitly links the Source and Target papers using only their combined QA pairs. The question must reference shared variables, mechanisms, or outcomes appearing in the combined QA, avoiding abstract or generic formulations. The answer is required to be logically entailed by the union of the two combined QA pairs and must not introduce external facts, unstated assumptions, or auxiliary evidence.

Step 3. Merge Complete QA The results from the previous step are merged into a single, coherent multi-hop question that (i) requires locating the Target Paper and (ii) involves reasoning across both papers. The LLM generates this merged question while preserving the intent of both components without introducing new information. The final answer is either identical to, or deterministically derived from, the inter-document answer, ensuring consistency and a well-structured reasoning chain.

Step 4. QA Validation (Optional) Each Complete QA is evaluated on the criteria in Table 4. Criterion is judged as either *Accept* or *Reject*; the *Overall Decision* is determined from these binary ratings. Items are retained only when the *Overall Decision* is *Accept*. In the validation experiment, no items were rejected during the final evaluation phase. To verify the functionality of the QA validation process, an additional analysis was conducted excluding the pre-rejected samples. Results showed that approximately **70% of pre-rejections originated from the Relational Appropriateness or Cross-reference Necessity** criteria, indicating that most filtering occurred when inter-paper connections were conceptually weak or unnecessary. This confirms that the validation pipeline effectively distinguishes between coherent and incoherent inter-document QA constructions.

Quality Assurance Four graduate-level annotators independently rate 150 randomly sampled Complete QA pairs across four criteria as shown table 4 using a 5-point Likert scale. Overall, the dataset showed strong linguistic and conceptual quality, with high mean scores and consistently coherent samples as shown in Figure 5.

4.5. Verifications

Evidence Coverage Validation Beyond pair-level validation, we examined whether the SHQA list faithfully represents section content and whether retrieval questions remain uniquely specified within each paper cluster. We first examined whether all salient information in a section was actually extracted as queries, since any unextracted information could undermine the uniqueness of the retrieval

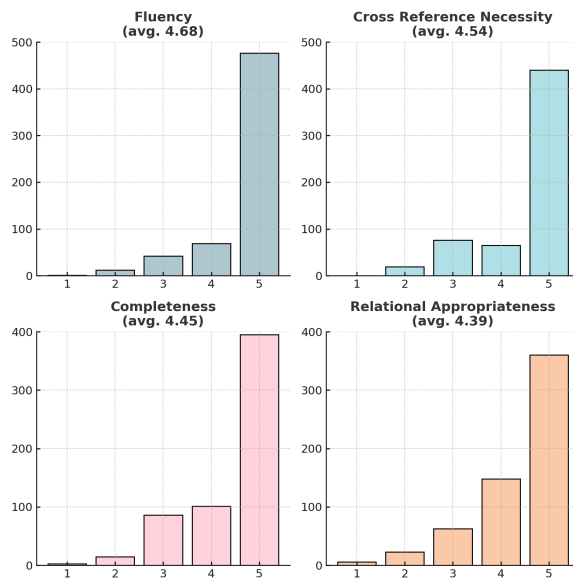


Figure 5: Results of multi-dimensional MHQA quality validation

Table 4: Criteria of MHQA quality validation.

Criterion	Description
Fluency	Whether the Complete QA is grammatically correct and naturally written.
Completeness	Whether the Complete QA can be answered solely based on the provided QA pair.
Cross-reference Necessity	Whether both papers are essential to answer the Complete QA.
Relational Appropriateness	Whether the two papers are conceptually suitable to be combined into a single QA.
Overall Decision	Synthesizes the above four criteria to make the final decision.

question. To justify using SHQA as a surrogate for section information, we conducted a post-hoc audit of 40 sections containing 202 SHQA items that had passed manual validation and checked whether the evidence texts covered the meaningful parts of each section. On average, 72% of sentences in a section appeared as SHQA evidence, and this coverage rose to 88% for sections such as abstract, result or conclusion, and discussion where independent facts are densely expressed. When a single evidence span conveyed multiple individual facts, the content was consistently split into multiple SHQA items so that each item captured one fact per QA. Notably, as shown in Table 5, a single evidence passage often yields multiple question-answer pairs, resulting in more than 12,000 multi-QA evidence units. This design enriches contextual diversity and contributes to building a higher-quality SHQA dataset.

Retrieval Question Validation We then assessed the uniqueness of the retrieval question within the paper cluster and manually audited all

test set retrieval questions. These findings support the use of SHQA as a faithful surrogate for section content and establish that retrieval questions uniquely identify their intended targets under our clustering and question design constraints.

Table 5: Example of Multi-QA Evidence

Evidence (E)	For instance, the U.S. CDC N1 assay had a recall of 99.0% and the U.S. CDC N2 assay had a 99.2% recall. PMC10021650
Meta	
Q1	What was the recall rate for the U.S. CDC N1 assay?
A1	The recall rate for the U.S. CDC N1 assay was 99.0%.
Q2	What was the recall rate for the U.S. CDC N2 assay?
A2	The recall rate for the U.S. CDC N2 assay was 99.2%.

5. Experiments

The reliability and suitability of AIM-SciQA were validated in previous sections. However, assessing the dataset against real-world task criteria represents a distinct challenge. Therefore, to examine its practical applicability, we conduct a two-stage evaluation—retrieval and QA—on the IM-SciQA dataset.

5.1. Retrieval Evaluation

We investigate the retrieval process for identifying the *Target Paper* using diverse retrievers. As described in Section 4, each query in IM-SciQA is associated with a predefined *paper cluster* constructed during the AIM-SciQA pipeline. A paper cluster is a compact set of 30 papers that includes one gold target and 29 distractor papers, derived either from citation-based or QA-based relational matching. This cluster serves as a constrained retrieval environment that closely mirrors realistic scientific search scenarios, where relevant literature typically resides within a semantically or citation-linked subset of the corpus.

To analyze the impact of such clustering constraints, we design three experimental environments: (1) *Paper Cluster* — retrieval is performed within the original cluster defined by the AIM-SciQA construction process; (2) *Random Cluster* — the same number of documents (30) is randomly sampled per query to remove relational coherence; and (3) *w/o Cluster* — retrieval is conducted over the full corpus of 8,211 papers. To ensure precise retrieval, we use the *Retrieval Question* only, rather than the Combined QA, as the query input.

Models We evaluate various embedding models to analyze differences in representation quality and retrieval accuracy. All models were implemented based on the LlamaIndex framework (Liu, 2022), and all experiments were conducted on an NVIDIA A100 (80GB) GPU.

Metrics To evaluate retrieval performance, we report **Hit@1, 3**, and **MRR@5** under the paper cluster setting ($N = 30$), and **Hit@1, 50**, and **MRR@50** under the full paper setting ($N = 8211$). **Hit@K** measures the proportion of queries for which the correct document appears within the top K retrieved results, while **MRR@K (Mean Reciprocal Rank)** averages the reciprocal rank of the first correct document, reflecting the top-k ranking quality.

5.2. IM-QA Evaluation

We perform the **IM-QA** process using the given Source Paper and the retrieved Target Paper with various LLMs. To disentangle the impact of the retrieval stage from the intrinsic IM-QA capability, we conduct evaluations under two settings: **Realistic** and **Oracle**. In the Realistic Setting, the top-1 document from the retrieval stage is provided as the Target Paper, whereas in the Oracle Setting, the gold Target Paper is supplied regardless of the retrieval outcome.

Models To obtain a representative view across model families, providers, and scales, we evaluate both open-source and proprietary LLMs. Inference for open-source models was performed using *vLLM* (Kwon et al., 2023) on an NVIDIA H100 GPU. To further compare against human-level performance, two biomedical domain experts were asked to answer 100 QA items under the Oracle Setting.

Metrics As the main metric, we adopt LLM-based Accuracy, following M3-SciQA. Specifically, a strong LLM (GPT-5) is prompted with the Question, Answer, and Prediction to evaluate whether the Answer and Prediction provide the same response to the Question, scored on a 0–0.5–1 scale. Traditional metrics such as token-level F1-score, ROUGE-L, and BERTScore (He et al., 2021) are also employed.

5.3. Results

Retrieval Performance for Various Cluster Settings The results of our retrieval evaluation are presented in Table 6. Under the Paper Cluster setting, lower-capacity models showed performance comparable to the Random Cluster setting, whereas higher-capacity models exhibited significantly lower scores in both Hit@1,3. A similar trend was observed in the Full Papers setting, where performance decreased despite the substantially larger search space (e.g., Hit@1 of 0.673 on 30 docs vs. 0.240 on 8,211 docs). These results indicate that **(1) the proposed Paper Clustering method effectively constructs a challenging and realistic candidate set that evaluates the fine-grained discriminative ability of**

Table 6: Retrieval performance across different settings. **Bold** indicates the best performance, and underline denotes the second best.

Model	Paper Cluster (N=30)			Random Cluster (N=30)			Full Papers (N=8211)		
	Hit@1	Hit@3	MRR@5	Hit@1	Hit@3	MRR@5	Hit@1	Hit@50	MRR@50
GritLM-7B (Muennighoff et al., 2024)	0.220	0.353	0.306	0.177	0.303	0.248	0.010	0.130	0.023
Jina-v2-base-en (Günther et al., 2023)	0.197	0.350	0.280	0.183	0.340	0.266	0.020	0.117	0.030
e5-base-v2 (Wang et al., 2022)	0.273	0.493	0.395	0.307	0.537	0.428	0.017	0.157	0.028
e5-mistral-7b-instruct (Wang et al., 2023)	0.300	0.473	0.390	0.273	0.430	0.356	0.030	0.203	0.058
Qwen3-Embedding-8B (Zhang et al., 2025)	0.407	0.597	0.554	0.483	0.607	0.555	0.080	0.397	0.136
llm-embedder (Xiao et al., 2024)	0.497	0.683	0.666	0.577	0.753	0.665	0.130	0.480	0.190
bge-large-en-v1.5 (Chen et al., 2024)	0.580	0.750	0.680	0.637	0.773	0.712	0.213	0.550	0.279
text-embedding-ada-002 (Greene et al., 2022)	0.600	0.783	0.695	0.667	0.830	0.752	0.207	0.553	0.272
voyage-3-large (VoyageAI, 2025)	0.673	0.830	0.761	0.737	0.887	0.811	0.240	0.637	0.325

Table 7: Comparison of model performance under *Oracle* and *Realistic* settings. **Bold** indicates the best performance, and underline denotes the second best (based on Accuracy). Accuracy in Oracle setting is measured over 100 QA instances.

Model	Oracle Setting				Realistic Setting			
	Acc	F1	ROUGE	BERTScore	Acc	F1	ROUGE	BERTScore
Human Expert	0.975	0.324	0.319	0.447	–	–	–	–
gemma-3-1b-it (Team et al., 2025)	0.265	0.337	0.301	0.340	0.260	0.332	0.292	0.328
Llama-3.2-3B-it (Grattafiori et al., 2024)	0.57	0.421	0.403	0.435	0.515	0.399	0.380	0.403
gemma-3-4b-it (Team et al., 2025)	0.590	0.430	0.403	0.465	0.515	0.418	0.390	0.453
gemma-3-27b-it (Team et al., 2025)	0.750	<u>0.517</u>	<u>0.495</u>	<u>0.556</u>	0.670	0.483	0.461	0.519
Qwen3-30B-A3B-it (Yang et al., 2025)	0.760	0.505	0.486	0.539	0.685	0.477	0.453	0.504
DeepSeek-R1 (Guo et al., 2025)	0.820	0.150	0.232	0.129	0.710	0.143	0.220	0.118
gpt-oss-20b (Agarwal et al., 2025)	0.835	0.454	0.438	0.489	0.705	0.423	0.405	0.448
claude-opus-4-1 (Anthropic, 2025)	<u>0.855</u>	0.476	0.464	0.522	<u>0.730</u>	0.430	0.412	0.465
gpt-5 (OpenAI, 2025)	0.900	0.520	0.507	0.570	0.775	<u>0.474</u>	<u>0.453</u>	<u>0.512</u>

retrieval models. Meanwhile, an error analysis of the Random Cluster setting revealed that many “incorrect” predictions actually contained information that could be regarded as valid alternative answers. Together with the observation in Section 4.5, this finding supports that **(2) the cluster-level Unique Retrieval Question selection process effectively mitigates the risk of multiple valid answers.**

Performance Separation under Oracle Setting

As shown in Table 7, modern proprietary LLMs achieve near human-level performance under the *Oracle Setting*, whereas smaller models such as *Gemma-3-1B*, *Gemma-3-4B*, and *LLaMA-3B* exhibit notably lower accuracy even on relatively simple relational reasoning tasks. The performance gap is more pronounced in the LLM-based Accuracy (GPT-Score) than in traditional metrics (F1, ROUGE-L, BERTScore), suggesting that smaller models tend to generate linguistically plausible yet factually inconsistent answers. This pattern indicates that the required information can be effectively retrieved and reasoned over when the model has sufficient capacity, providing empirical evidence that both the question hints and the target information are well-grounded and meaningful.

Robustness and Sensitivity under Realistic Setting Meanwhile, as shown in Table 7, all models experience a performance degradation under the

Table 8: CIM-SciQA Results (Realistic Setting)

Model	Acc	F1	ROUGE	BERTScore
gemma-3-1b-it (Team et al., 2025)	0.446	0.330	0.295	0.366
gemma-3-4b-it (Team et al., 2025)	0.635	0.366	0.329	0.405
Llama-3.2-3B-it (Grattafiori et al., 2024)	0.608	0.367	0.339	0.379
gemma-3-27b-it (Team et al., 2025)	0.798	<u>0.427</u>	<u>0.402</u>	<u>0.449</u>
DeepSeek-R1 (Guo et al., 2025)	0.784	0.131	0.205	0.110
Qwen3-30B-A3B-it (Yang et al., 2025)	0.824	0.425	0.392	0.439
gpt-oss-20b (Agarwal et al., 2025)	<u>0.851</u>	0.373	0.346	0.381
gpt-5 (OpenAI, 2025)	0.926	0.432	0.411	0.478

Realistic Setting, where the retrieved document is used as the target. However, the degree of decline does not scale directly with the retrieval accuracy (Hit@1 = 0.673), implying that the reasoning process anchored on the Source Paper remains partially robust to retrieval noise. Interestingly, stronger models that perform well under the Oracle Setting tend to show larger drops when retrieval is imperfect, suggesting that more capable models rely more heavily on precise evidence alignment. This finding supports the view that providing the correct supporting document substantially facilitates inter-document reasoning and that the proposed dataset effectively evaluates such reasoning under realistic retrieval conditions.

5.4. CIM-SciQA Evaluation

We evaluated **CIM-SciQA**, a citation-guided variant of AIM-SciQA, under the same evaluation configuration as the IM-SciQA dataset, except that we

did not employ claude-opus-4-1 or human expert evaluation due to limited budget. In this setting, the retrieval task of CIM-SciQA is conducted by locating appropriate in-text citation symbols within the source paper, while the realistic QA task is prompted to generate answers by referencing the full text of both the source paper and the paper corresponding to the cited symbol. The evaluation results are presented in Table 8.

CIM-SciQA achieved performance comparable to the Oracle setting of IM-SciQA (Table 6), substantially surpassing the Realistic setting. Even when considering the inherent complexity of the documents included in both datasets, this gap represents a significant difference. We attribute this to the nature of citation-based questions, which are grounded in explicitly linked documents and therefore tend to form QA pairs of relatively lower difficulty.

6. Conclusion

We present **AIM-SciQA**, an automatic framework that constructs scientific IM-QA datasets, and introduce the resulting dataset, **IM-SciQA**. Our method systematically combines SHQA extraction through both semantic and citation-based matching. Using AIM-SciQA, we construct the IM-SciQA dataset from 8,211 PubMed Central articles, totaling 411k SHQA and 13K MHQA pairs with rigorous human and automatic validation. Our experimental analyses show that recent LLMs remain inferior to human experts in factual consistency and multi-hop reasoning, when they exhibit strong linguistic fluency, and the gap is most pronounced under realistic retrieval settings. These findings underscore the value of IM-SciQA as a challenging benchmark for evaluating retrieval-augmented reasoning and scientific comprehension in modern LLMs. We hope AIM-SciQA provides a durable foundation for research on scalable scientific QA generation, improved retrieval-reasoning integration, and trustworthy multi-hop inference.

7. Limitations

Our current framework does not incorporate multi-modal information such as figures or tables. This limitation arises due to lack of sufficiently powerful models capable of being effectively integrated into our automatic generation pipeline. In addition, to ensure reliable multi-hop QA (MHQA) performance during automatic generation, we employed high-capacity large language models (LLMs) for combining multiple QA instances. While this approach improves answer consistency and reasoning quality, it incurs substantial computational and financial costs. SHQA is designed to capture

salient, answerable facts, not every detail of a paper. Consequently, it does not fully replicate the entire document. Still, the combination of strict exact-evidence matching, similarity-difference filtering, high inter-rater reliability ($AC1 = 0.94$), and section-level coverage (72–88%) supports its adequacy as a pragmatic proxy for inter-document reasoning. Furthermore, our current work mainly focuses on inter-document reasoning, without exploring deeper multi-hop chains or diverse reasoning types. Expanding SHQA to include more complex reasoning hops and richer inference structures remains an important direction for future work.

8. Acknowledgements

This work was partly supported by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2020-0-01361, Artificial Intelligence Graduate School Program (Yonsei University)).

9. Bibliographical References

- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Anthropic. 2025. Claude opus 4.1. <https://www.anthropic.com/news/claude-opus-4-1>. News release.
- Abhinand Balachandran. 2024. [Medembed: Medical-focused embedding models](#).
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216*.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, et al. 2025. Towards an ai co-scientist. *arXiv preprint arXiv:2502.18864*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

- Ryan Greene, Ted Sanders, Lilian Weng, and Arvind Neelakantan. 2022. New and improved embedding model. <https://openai.com/index/new-and-improved-embedding-model>. OpenAI Blog.
- Michael Günther, Jackmin Ong, Isabelle Mohr, Alaeddine Abdessalem, Tanguy Abel, Mohammad Kalim Akram, Susana Guzman, Georgios Mastrapas, Saba Sturua, Bo Wang, et al. 2023. Jina embeddings 2: 8192-token general-purpose text embeddings for long documents. *arXiv preprint arXiv:2310.19923*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Kilem Li Gwet. 2008. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Seonjeong Hwang, Yunsu Kim, and Gary Geunbae Lee. 2024. Explainable multi-hop question generation: An end-to-end approach without intermediate question labeling. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6855–6866.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Seongyun Lee, Hyunjae Kim, and Jaewoo Kang. 2023. Liquid: A framework for list question answering dataset generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13014–13024.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Zefeng Lin, Weidong Chen, Yan Song, and Yongdong Zhang. 2024. Prompting few-shot multi-hop question generation via comprehending type-aware semantics. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3730–3740.
- Jerry Liu. 2022. [LlamaIndex](#).
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*.
- Niklas Muennighoff, SU Hongjin, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and Douwe Kiela. 2024. Generative representational instruction tuning. In *The Thirteenth International Conference on Learning Representations*.
- OpenAI. 2025. [Introducing gpt-5](https://openai.com/index/introducing-gpt-5). <https://openai.com/index/introducing-gpt-5>. Blog post.
- Zhengliang Shi, Weiwei Sun, Shen Gao, Pengjie Ren, Zhumin Chen, and Zhaochun Ren. 2024. Generate-then-ground in retrieval-augmented generation for multi-hop question answering. *arXiv preprint arXiv:2406.14891*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- VoyageAI. 2025. voyage-3-large: the new state-of-the-art general-purpose and multilingual embedding model that ranks first across eight evaluated domains spanning 100 datasets. <https://blog.voyageai.com/2025/01/07/voyage-3-large>. Blog post.
- David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hannaneh

- Hajishirzi. 2022. Scifact-open: Towards open-domain scientific claim verification. *arXiv preprint arXiv:2210.13777*.
- Yuwei Wan, Yixuan Liu, Aswathy Ajith, Clara Grazian, Bram Hoex, Wenjie Zhang, Chunyu Kit, Tong Xie, and Ian Foster. 2024. Sciqag: A framework for auto-generated science question answering dataset with fine-grained evaluation. *arXiv preprint arXiv:2405.09939*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxiang Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2023. Improving text embeddings with large language models. *arXiv preprint arXiv:2401.00368*.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. 2024. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pages 641–649.
- An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, et al. 2025. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.
- Dasigi, Pradeep and Lo, Kyle and Beltagy, Iz and Cohan, Arman and Smith, Noah A and Gardner, Matt. 2021. *A Dataset of Information-Seeking Questions and Answers Anchored in Research Papers*. Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.
- Jin, Qiao and Dhingra, Bhuwan and Liu, Zhengping and Cohen, William W and Lu, Xinghua. 2019. *Pubmedqa: A dataset for biomedical research question answering*. *arXiv preprint arXiv:1909.06146*.
- Koreeda, Yuta and Manning, Christopher D. 2021. *ContractNLI: A Dataset for Document-level Natural Language Inference for Contracts*. Findings of the Association for Computational Linguistics: EMNLP 2021.
- Krithara, Anastasia and Nentidis, Anastasios and Bougiatiotis, Konstantinos and Paliouras, Georgios. 2023. *BioASQ-QA: A manually curated corpus for Biomedical Question Answering*. Nature Scientific Data.
- Lee, Yookyung and others. 2023. *QASA: Advanced Question Answering on Scientific Articles*. PMLR.
- Li, Chuhan and Shangquan, Ziyao and Zhao, Yilun and Li, Deyuan and Liu, Yixin and Cohan, Arman. 2024. *M3SciQA: A multi-modal multi-document scientific QA benchmark for evaluating foundation models*. *arXiv preprint arXiv:2411.04075*.
- Lu, Peng and Bansal, Mohit and Xia, Tony and Liu, Siyuan and Sun, Shiyu and others. 2022. *Learn to Explain: Multimodal Reasoning via Thought Chains for Science Question Answering*. Advances in Neural Information Processing Systems.
- Möller, Timo and Reina, Anthony and Jayakumar, Raghavan and Pietsch, Malte. 2020. *COVID-QA: A Question Answering Dataset for COVID-19*. Association for Computational Linguistics. Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020.
- Park, Jiwon and Pyeon, Seohyun and Kim, Jinwoo and Cabal, Rina Carines and Ding, Yihao and Han, Soyeon Caren. 2025. *DocHop-QA: Towards Multi-Hop Reasoning over Multimodal Document Collections*. *arXiv preprint arXiv:2508.15851*.
- Pramanick, Shraman and Chellappa, Rama and Venugopalan, Subhashini. 2024. *Spiqa: A dataset for multimodal question answering on scientific papers*. Advances in Neural Information Processing Systems.

10. Language Resource References

- Amiri, Mahmoud and Bocklitz, Thomas. 2025. *ChemRxivQuest: A Curated Chemistry Question-Answer Database Extracted from ChemRxiv Preprints*. *arXiv preprint arXiv:2505.05232*.
- Baumgrtner, Tim and Briscoe, Ted and Gurevych, Iryna. 2025. *PeerQA: A Scientific Question Answering Dataset from Peer Reviews*. Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers).

A. Prompts

A.1. Single-hop QA Generation Prompt

You are an expert in reading comprehension.
You will be provided with a single paragraph from a scientific paper.
Read the paragraph carefully and identify meaningful information (entities, facts, relations, events) that can be directly and unambiguously extracted from it.

For each identified meaningful fact, you must generate a JSON object with the following three fields:

question: Generate one clear and specific question. The question must detail the context or conditions (e.g., the specific patient group, timeframe, or situation) mentioned in the text to pinpoint the fact accurately.

evidence: Extract the minimal contiguous text span from the paragraph that provides the evidence for the answer. This must be an exact quote from the text.

answer: Based on the extracted evidence, formulate a natural complete sentence that directly answers the question.

Rules:

Do not use any outside knowledge or inference.
All information must be sourced only from the given paragraph.

The evidence field must be an exact extraction without modification.

Avoid generating reasoning-type (multi-hop) questions.

Return the result as a JSON list with the following structure:

```
{  
  "question": "string",  
  "evidence": "string",  
  "answer": "string"  
}
```

A.2. Multi-hop QA Generation Prompt

A.2.1. Reasoning Prompt

You are a scientific QA generator for inter-document multi-hop question construction.

Your task is to produce a structured QA reasoning process that constructs an inter-document multi-hop question across multiple scientific papers.

Step 0. Validation (not included in output)

Before generating any reasoning output, verify whether the Source QA and Target QA form a meaningful multi-hop relationship such as comparison, causation, conceptual linkage, or inference.

If no logical or conceptual overlap exists, output only:

Cannot generate inter-document multi-hop question.

Reasoning Components

Find Target Paper

Transform the Unique Target QA question into a question that asks which paper addresses that topic.

The answer must identify the target paper using the format:

Identified target paper: '[Target Paper Title]'.

Generate Inter-document QA

Construct a comparative or integrative question connecting the Source Paper and Target Paper using the Core QA pair.

The answer must be logically derived only from the provided QA pair.

Merge Complete QA

Combine the Find Target Paper question and the Inter-document QA question into a single multi-hop question.

The final answer must be identical to or directly derived from the Inter-document QA answer.

QA Validation

Evaluate the generated QA using the following criteria:

Fluency

Completeness

Cross-reference Necessity

Relational Appropriateness

Each criterion is labeled Accept or Reject and followed by an overall Decision.

Output Format

<outputs>

<component type="Find Target Paper">

<question>...</question>

<answer>...</answer>

</component>

<component type="Generate Inter-document QA">

<question>...</question>

<answer>...</answer>

</component>

<component type="Merge Complete QA">

<question>...</question>

<answer>...</answer>

```
</component>

<component type="QA Validation">
  <score type="Fluency">...</score>
  <score type="Completeness">...</score>
  <score type="Cross-reference Necessity">...</score>
  <score type="Relational Appropriateness">...</score>
  <score type="Decision">...</score>
</component>
</outputs>
```

A.2.2. Input Template

Now Your Turn

```
<inputs>
  <source_paper>
    <title>{source_paper_title}</title>
    <section_name>{source_section_name}</section_name>
    <core_qa>
      <question>{core_source_q}</question>
      <answer>{core_source_a}</answer>
    </core_qa>
  </source_paper>
  <target_paper>
    <title>{target_paper_title}</title>
    <section_name>{target_section_name}</section_name>
    <core_qa>
      <question>{core_target_q}</question>
      <answer>{core_target_a}</answer>
    </core_qa>
    <unique_qa>
      <question>{unique_target_q}</question>
      <answer>{unique_target_a}</answer>
    </unique_qa>
  </target_paper>
</inputs>
```